

# Dell EMC™ Unisphere for PowerMax™

Version 9.0.0

Online Help (PDF version)

Copyright © 2012-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published May 2018

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

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

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

Dell EMC  
Hopkinton, Massachusetts 01748-9103  
1-508-435-1000 In North America 1-866-464-7381  
[www.DellEMC.com](http://www.DellEMC.com)

# CONTENTS

<b>Tables</b>		<b>21</b>
<b>Chapter 1</b>	<b>Introduction</b>	<b>25</b>
	Unisphere Online Help.....	26
	Capacity information.....	29
<b>Chapter 2</b>	<b>Getting Started</b>	<b>31</b>
	Operating as the initial setup user.....	32
	Viewing Home Dashboard view - All Storage Systems.....	32
	Viewing Home Dashboard view - Specific Storage System.....	34
	Viewing system performance view.....	34
	Viewing the System Health Dashboard.....	35
	Understanding the system health score.....	37
	Viewing Storage Group Compliance view.....	39
	Viewing Capacity dashboard view.....	40
	Viewing Replication dashboard.....	41
	Discovering storage systems.....	43
	Refreshing storage system information.....	44
	Viewing product version information.....	44
	Searching for storage objects.....	45
	Modifying server logging levels.....	46
	Exiting the console.....	46
	Getting help.....	46
<b>Chapter 3</b>	<b>Administration</b>	<b>47</b>
	Managing settings.....	48
	Setting preferences.....	49
	Backing up the database server.....	50
	Viewing database backups.....	51
	Deleting database backups.....	51
	Alert settings.....	51
	Alerts.....	52
	Alert policies.....	56
	Threshold alerts.....	56
	Configuring email notifications.....	59
	Editing subscriptions .....	60
	Performance thresholds and alerts.....	60
	Service level alert policies.....	63
	Server alerts.....	66
	Security.....	67
	Authentication.....	67
	Understanding user authorization.....	70
	View Certificate dialog box.....	76
	Local Users.....	76
	Viewing user sessions.....	80
	Roles and associated permissions.....	80
	Link and launch.....	83
	Creating link-and-launch client registrations.....	83

Editing link-and-launch client registrations.....	84
Deleting link-and-launch client registrations.....	84
Viewing link and launch client registrations.....	84
Managing Database Storage Analyzer (DSA) environment preferences.....	85
Managing data protection preferences.....	85
Viewing authentication authority information.....	86
Local User and Authorization operations.....	87
Link and Launch operations.....	87
Entering PIN number .....	87
Report operations.....	87

<b>Chapter 4</b>	<b>Storage Management</b>	<b>89</b>
	Understanding Storage Management.....	90
	Tag and Untag operations.....	91
	Viewing Storage Group Demand Reports.....	91
	Viewing Service Level Demand Reports.....	92
	Viewing CKD volumes.....	92
	Viewing CKD volumes in CU image.....	93
	Viewing Storage Group Compliance view.....	94
	Dialog displayed when there is less than one week's data collected.....	96
	Setting volume emulation.....	96
	FAST association operations.....	97
	Removing DATA volumes.....	97
	Mapping volume operations.....	97
	Rename operations.....	98
	Provisioning storage .....	98
	Using the Provision Storage wizard.....	100
	Provisioning storage for mainframe.....	104
	Using the Provision Storage wizard for mainframe.....	104
	Provisioning storage.....	107
	Using the Provision Storage wizard.....	108
	Suitability Check restrictions.....	111
	Creating storage groups.....	112
	Adding volumes to storage groups.....	114
	Copying volumes between storage groups.....	114
	Moving volumes between storage groups.....	115
	Removing volumes from storage groups.....	116
	Storage Group operations.....	116
	Expanding storage groups.....	116
	Expanding ProtectPoint storage groups.....	118
	Modifying storage groups.....	119
	Renaming storage groups.....	121
	Protecting storage groups.....	122
	Converting storage groups to cascaded.....	129
	Changing Storage Resource Pools for storage groups.....	129
	Adding or removing cascaded storage groups.....	130
	Renaming storage groups.....	131
	Deleting storage groups.....	131
	Setting host I/O limits.....	132
	Splitting storage groups.....	133
	Merging storage groups.....	134
	Managing VP compression on thin volumes in storage groups.....	134
	Viewing storage groups.....	135
	Select Storage Resource Pool.....	148
	Select SSID.....	149

Task in Progress.....	149
Select SRDF group.....	149
Editing storage group volume details .....	149
Editing storage group details .....	150
Modify Custom Capacity dialog box.....	150
Understanding FAST .....	151
Understanding service levels.....	151
Understanding Storage Resource Pool details.....	153
Service level compliance.....	157
Symmetrix tiers.....	161
FAST policies.....	166
Pinning and unpinning volumes.....	173
Time windows.....	174
Understanding Workload Planner.....	176
Managing volumes.....	177
Managing volumes .....	178
Creating volumes.....	178
Deleting volumes.....	188
Duplicating volumes.....	188
Assigning array priority to individual volumes.....	189
Assigning array priority to groups of volumes.....	189
Changing volume configuration.....	190
Expanding existing volumes.....	191
Mapping volumes.....	192
Unmapping volumes.....	193
Setting optimized read miss.....	193
Setting volume status.....	194
Setting volume attributes.....	195
Setting volume identifiers.....	196
Setting volume names.....	196
Setting copy pace (QoS) for device groups.....	197
QOS for replication.....	197
Setting copy pace (QoS) for storage groups.....	197
Setting copy pace (QoS) for volumes.....	198
Managing Meta Volumes.....	199
Viewing CKD volumes.....	203
Viewing CKD volume details.....	204
Viewing CKD volume front end paths.....	207
Viewing DLDEV volumes.....	207
Viewing DLDEV volume details.....	207
Viewing masking information.....	208
Viewing meta volumes.....	208
Viewing meta volume details.....	209
Viewing meta volume meta members.....	212
Viewing meta volume member details.....	212
Viewing other pool information.....	214
Viewing private volumes .....	215
Viewing private volume details.....	215
Viewing regular volumes.....	217
Viewing regular volume details.....	218
Viewing reserved volumes.....	221
Viewing reserved volume details.....	221
Viewing SAVE volumes .....	222
Viewing SAVE volume details.....	222
Viewing storage resource pool information.....	223
Viewing thin volumes .....	223

Viewing thin volume details.....	223
Viewing thin volume bound pool information.....	223
Viewing virtual volumes.....	224
Viewing virtual volume details.....	225
Viewing volume back end paths.....	227
Viewing volume FBA front end paths.....	228
Viewing volume RDF information.....	228
Select Volume Range dialog box.....	229
Advanced Options dialog.....	230
Viewing disk groups.....	230
Viewing disk group details.....	230
Viewing disks in disk group.....	231
Viewing disk details.....	231
Viewing disk hyper volumes.....	232
Viewing hyper volume details.....	233
Viewing volumes for disk.....	234
Viewing paths for disks.....	234
Viewing spare disks in disk group.....	235
Viewing spare disk details.....	235
Removing disks from disk groups.....	236
Deleting disk groups.....	237
Renaming disk groups.....	237
Creating DATA volumes.....	237
Activating and deactivating DATA volumes.....	238
Enabling and disabling DATA volumes.....	238
Start and stop draining DATA volumes.....	239
Viewing DATA volumes.....	239
Viewing DATA volume details.....	239
Creating thin pools.....	240
Expanding thin pools.....	241
Draining thin pools.....	241
Starting and stopping thin pool write balancing.....	242
Deleting thin pools.....	243
Adding or removing thin pool members.....	243
Enabling and disabling thin pool members.....	244
Managing thin pool allocations.....	244
Viewing thin pools.....	245
Viewing thin pool details.....	246
Viewing bound volumes for a thin pool.....	247
Viewing DATA volumes for a thin pool.....	250
Viewing details on DATA volumes in thin pools.....	250
Viewing other volumes for thin pools.....	252
Managing thin pool capacity.....	253
Allocate/Free/Reclaim dialogs.....	254
Creating or Expanding or Modifying thin pools.....	255
Creating thin volumes.....	256
Binding/Unbinding/Rebinding thin volumes.....	257
Understanding Virtual LUN Migration.....	258
Viewing VLUN migration sessions.....	259
Viewing VLUN migration session details.....	259
Terminating a VLUN migration session.....	260
VLUN Migration dialog box.....	260
Select VLUN Migration Session Target dialog box.....	260
Migrating regular storage group volumes.....	261
Migrating regular volumes.....	261
Migrating thin storage group volumes.....	262

Migrating thin volumes.....	262
Understanding Federated Tiered Storage .....	263
Viewing external storage.....	263
Virtualizing external LUNs.....	264
Virtualizing external LUNs.....	265
Removing external LUNs.....	266
Understanding storage templates.....	267
Creating storage templates.....	267
Viewing storage templates.....	270
Modifying storage templates.....	271
Deleting storage templates.....	272
Understanding FAST.X.....	272
Viewing external disks.....	272
Adding external disks.....	274
Removing external disks or External LUNs.....	274
Working with external disks.....	275
Start draining external disks.....	275
Stop draining external disks.....	276
Activating external disks.....	276
Viewing reservations.....	277
Viewing reservation details.....	277
Releasing reservations.....	278
Managing vVol.....	278
Viewing storage containers.....	279
Viewing storage container details.....	280
Creating storage containers.....	281
Modifying storage containers.....	282
Deleting storage containers.....	282
Viewing storage resources.....	282
Viewing storage resource details.....	283
Viewing storage resource related SRPs.....	283
Adding storage resources to storage containers.....	284
Modifying storage resources.....	285
Removing storage resources from storage containers.....	285
Viewing protocol endpoints.....	285
Viewing protocol endpoint details.....	286
Provisioning protocol endpoints to hosts.....	286
Deleting protocol endpoints.....	287
Configuring the VASA provider connection.....	287
Understanding compression.....	287
Viewing the SRP efficiency details .....	288
Viewing compressibility reports.....	288
Viewing a storage group's compression ratio.....	289
Viewing a volume's compression details.....	289
Viewing compression status using the VVol Dashboard.....	290
Viewing the compression efficiency dashboard.....	290
 <b>Chapter 5</b>	
<b>Host Management</b>	<b>291</b>
Understanding Host Management.....	292
Creating hosts.....	292
Adding initiators to hosts.....	293
Adding initiator to host.....	294
Removing initiators from hosts.....	294
Modifying hosts.....	294
Renaming hosts/host groups.....	295

Setting host or host group port flags.....	296
Deleting hosts/host groups.....	296
Viewing hosts/host groups.....	296
Viewing host/host group details.....	297
Viewing host initiators.....	298
Host/Host group flags.....	299
Host I/O limits dialog box.....	300
Host Group filtering rules.....	300
Select Storage Resource Pool.....	301
Provisioning storage.....	302
Creating host groups.....	302
Adding hosts to host groups.....	303
Removing hosts from host groups.....	303
Modifying host groups.....	304
Renaming hosts/host groups.....	304
Setting host or host group port flags.....	305
Deleting hosts/host groups.....	305
Viewing hosts/host groups.....	306
Viewing host/host group details.....	306
Creating masking views.....	307
Renaming masking views.....	308
Deleting masking views.....	308
Viewing masking views.....	308
Viewing masking view connections.....	309
Viewing masking view details.....	311
Set Dynamic LUN Addresses.....	311
Setting initiator port flags.....	311
Setting initiator attributes.....	312
Renaming initiator aliases.....	312
Replacing initiators.....	313
Removing masking entries.....	313
Viewing initiators.....	313
Viewing initiator details.....	314
Viewing volumes associated with host initiator.....	315
Viewing details of a volume associated with initiator.....	316
Creating port groups.....	316
Deleting port groups.....	317
Adding ports to port groups.....	317
Removing ports from port groups.....	318
Renaming port groups.....	319
Viewing port groups.....	319
Viewing port groups details.....	319
Viewing ports in port group.....	320
Viewing port details.....	321
Volume Set Addressing.....	321
Viewing host IO limits.....	321
Managing storage for Mainframe.....	322
Provisioning storage for mainframe.....	324
Using the Provision Storage wizard for mainframe.....	324
Viewing splits.....	327
Viewing CU images.....	328
Viewing CU image details.....	329
Creating CKD volumes.....	330
Editing CKD volume capacities.....	331
Expanding CKD volumes.....	331
z/OS map from the CU image list view.....	332



	z/OS unmap from the CU image list view.....	333
	z/OS map from the volume list view .....	333
	z/OS unmap from the volume list view.....	334
	z/OS map from the Volumes (Storage Groups) list view.....	335
	z/OS unmap from the Volumes (Storage Groups) list view.....	335
	Adding an alias range to a CU image.....	336
	Removing an alias range from a CU image .....	337
	Setting the base address.....	337
	Understanding All Flash Mixed FBA/CKD support .....	337
	Mapping CKD volumes.....	340
	Unmapping CKD volumes.....	341
	Copying CU image mapping.....	342
	Available Volume for EA/EF Mapping dialog box.....	342
	Base Addresses in Use dialog box.....	342
	Select SSID dialog box.....	342
	Viewing CKD volumes in CU image.....	342
	Creating PowerPath hosts.....	343
	Viewing PowerPath hosts.....	344
	Viewing PowerPath hosts details.....	344
	Viewing PowerPath Host Virtual Machines.....	345
	Viewing host cache adapters.....	346
<b>Chapter 6</b>	<b>Data Protection</b>	<b>347</b>
	Understanding Data Protection Management.....	348
	Creating device groups.....	348
	Adding volumes to device groups.....	349
	Removing volumes from device groups.....	350
	Setting consistency protection.....	350
	Renaming device groups.....	351
	Deleting device groups.....	351
	Viewing device groups.....	351
	Viewing device group details.....	352
	Viewing volumes in device group.....	353
	Understanding TimeFinder/Clone operations.....	354
	Managing TimeFinder/Clone sessions.....	354
	Creating clone copy sessions.....	355
	Activating clone copy sessions.....	357
	Recreating clone copy sessions.....	358
	Creating clone snapshots.....	359
	Modifying clone copy sessions.....	360
	Restoring data from target volumes.....	361
	Splitting clone volume pairs.....	362
	Terminating clone copy sessions.....	363
	Viewing clone pairs.....	364
	Viewing clone pair details.....	365
	Clone copy session options.....	365
	Understanding TimeFinder/Snap operations.....	367
	Managing TimeFinder/Snap sessions.....	368
	Creating virtual copy sessions.....	369
	Activating virtual copy sessions.....	370
	Creating snapshots.....	371
	Duplicating virtual copy sessions.....	372
	Recreating virtual copy sessions.....	373
	Restoring virtual copy sessions.....	374
	Terminating virtual copy sessions.....	375

Viewing snap pair details.....	376
Viewing snap pairs.....	376
Snap session options.....	377
Set TimeFinder Snap Pairs dialog box.....	378
Managing TimeFinder/Mirror sessions.....	378
Creating Snapshots.....	379
Restoring BCV pairs.....	380
Splitting BCV pairs.....	381
Cancelling BCV pairs.....	381
Viewing mirror pairs.....	382
Viewing mirror pair details.....	382
TimeFinder/Mirror session options.....	383
Setting TimeFinder/Mirror pairs.....	385
Managing TimeFinder SnapVX.....	385
Creating snapshots.....	387
Modifying TimeFinder SnapVX snapshots.....	389
Linking to snapshots.....	390
Relinking to snapshots.....	391
Unlinking from snapshots.....	392
Restoring snapshots.....	393
Setting snapshots to automatically terminate.....	393
Setting "Secure" status on an existing snapshot.....	394
Terminating snapshots.....	395
Setting copy mode for snapshots.....	396
Viewing snapshots.....	396
Viewing snapshot details.....	397
Viewing snapshot links.....	398
Viewing snapshot link details.....	399
Viewing snapshot source volumes.....	399
Viewing snapshot source volume details.....	400
Viewing snapshot source volume linked volumes.....	401
RBAC roles for performing local and remote replication actions..	402
Managing remote replication sessions.....	402
Creating SRDF connections.....	403
Creating SRDF pairs.....	404
Deleting SRDF pairs.....	407
Moving SRDF pairs.....	408
Setting SRDF mode.....	409
Viewing SRDF volume pairs.....	410
Viewing SRDF volume pair details.....	412
Viewing SRDF volume pair details.....	414
Viewing SRDF protected storage group pairs .....	415
Viewing SRDF protected storage group pair properties.....	417
Deleting SRDF pairs.....	419
Deleting SRDF pairs from the SRDF List Volumes View.....	421
Establishing SRDF pairs.....	421
Failing over.....	422
Failing back.....	423
Invalidating R1/R2 volumes.....	424
Making R1/R2 volumes ready.....	425
Making R1/R2 volumes not ready.....	426
Read/write disabling R2 volumes.....	427
Read/write enabling R1/R2 volumes.....	428
Resuming SRDF links.....	429
Read/write disabling R1/R2 volumes.....	429
Refreshing R1 or R2 volumes.....	430

Setting SRDF/A controls to prevent cache overflow.....	431
Setting consistency protection.....	432
Resetting original device identity.....	432
Restoring SRDF pairs.....	433
Setting bias location.....	434
Setting the SRDF GCM flag.....	434
Setting volume status.....	435
Splitting SRDF pairs.....	436
Suspending SRDF pairs.....	436
Swapping SRDF personalities.....	438
Updating R1 volumes.....	438
SRDF session options.....	439
SRDF session modes.....	442
RBAC roles for performing local and remote replication actions..	443
Understanding Virtual Witness.....	444
Adding SRDF Virtual Witness instances.....	445
Removing SRDF Virtual Witness instances.....	445
Set state for SRDF Virtual Witness instances.....	446
Viewing SRDF Virtual Witness instances.....	447
Viewing SRDF Virtual Witnesses details.....	447
Creating SRDF/A DSE pools.....	448
Deleting SRDF/A DSE pools.....	448
Adding volumes to SRDF/A DSE pools.....	449
Removing volumes from SRDF/A DSE pools.....	449
Enabling all volumes in SRDF/A DSE pools.....	449
Disabling all volumes in SRDF/A DSE pools.....	449
Viewing SRDF/A DSE pools.....	450
Viewing SRDF DSE pool details.....	450
Creating TimeFinder/Snap pools.....	451
Adding volumes to TimeFinder/Snap pools.....	452
Enabling all volumes in TimeFinder/Snap pools.....	452
Disabling all volumes in TimeFinder/Snap pools.....	452
Deleting TimeFinder/Snap Pools.....	452
Removing volumes from TimeFinder/Snap pools.....	453
Viewing TimeFinder/Snap pools.....	453
Viewing TimeFinder/Snap pool details.....	454
Viewing SRDF group volumes.....	454
Viewing SRDF protected storage groups.....	455
Viewing related SRDF groups.....	457
Creating SRDF groups.....	457
Modifying SRDF groups.....	459
Setting SRDF/A DSE attributes.....	461
Setting SRDF/A group attributes.....	461
Setting SRDF/A pace attributes.....	462
Swapping SRDF groups.....	463
Setting consistency protection.....	463
Deleting SRDF groups.....	464
Viewing SRDF groups.....	464
Viewing SRDF group details.....	465
Viewing SRDF protected device groups.....	467
Resuming SRDF links.....	468
Viewing SRDF group volumes.....	469
SRDF/A control actions.....	469
RDFA flags.....	470
SRDF group modes.....	471
Understanding RecoverPoint.....	471

Tagging and untagging volumes for RecoverPoint (storage group level).....	472
Tagging and untagging volumes for RecoverPoint (volume level)....	472
Untagging RecoverPoint tagged volumes.....	473
Viewing RecoverPoint copies.....	473
Viewing RecoverPoint copy details.....	474
Viewing RecoverPoint sessions.....	475
Viewing RecoverPoint session details.....	475
Viewing RecoverPoint storage groups.....	476
Viewing RecoverPoint tagged volumes.....	476
Viewing RecoverPoint tagged volume details.....	477
Protecting storage groups using RecoverPoint.....	479
Viewing RecoverPoint volumes.....	480
Viewing RecoverPoint clusters.....	481
Viewing RecoverPoint cluster details.....	482
Viewing RecoverPoint splitters.....	482
Viewing RecoverPoint appliances.....	483
RecoverPoint systems.....	483
RecoverPoint consistency groups.....	486
RecoverPoint replication sets.....	488
RecoverPoint links.....	489
Creating Open Replicator copy sessions.....	490
Activating Open Replicator session.....	491
Recreating Open Replicator sessions.....	492
Restoring Open Replicator sessions.....	492
Renaming Open Replicator sessions.....	492
Removing Open Replicator sessions.....	493
Setting Open Replicator session background copy mode.....	493
Setting Open Replicator session donor update off.....	493
Setting Open Replicator session front end zero detection off....	493
Setting Open Replicator session pace.....	494
Setting Open Replicator ceiling.....	494
Terminating Open Replicator sessions.....	494
Viewing Open Replicator sessions.....	495
Viewing Open Replicator session details.....	496
Viewing Open Replicator SAN View.....	496
Open Replicator session options.....	497
Open Replicator flags.....	499
Understanding non-disruptive migration (NDM).....	500
Preparing a non-disruptive migration (NDM) session.....	500
Creating a non-disruptive migration (NDM) session.....	502
Viewing the non-disruptive migration (NDM) sessions list.....	504
Viewing migration details.....	504
Readying the migration target.....	506
Cutting over a migration session.....	506
Synchronizing data after non-disruptive migration (NDM) cutover ..	507
Committing a migration session.....	507
Cancelling a migration session.....	508
Recovering a migration session.....	509
Viewing migration environments.....	509
Setting up a migration environment.....	510
Removing a migration environment.....	510
Viewing the authorized users and groups details.....	510
Expanding remote volumes .....	511

	Setting a device identity.....	511
	Editing storage group volume details .....	512
	Editing storage group details .....	513
	Replication state severities.....	513
	Managing space reclamation.....	514
	Advanced Options dialog.....	515
<b>Chapter 7</b>	<b>Performance Management</b>	<b>517</b>
	Understanding Performance Management.....	518
	Performance Dashboards.....	518
	Viewing dashboards.....	519
	Using default dashboards.....	520
	Using the All Arrays overview dashboard.....	521
	Creating a dashboard with charts.....	522
	Editing a template dashboard.....	523
	Copying a dashboard.....	523
	Editing dashboards.....	523
	Deleting dashboards.....	523
	Running a report from the dashboard.....	524
	Saving a dashboard as a template.....	524
	Saving dashboard changes.....	524
	Scheduling a report from the dashboard.....	525
	Navigating to the Details view.....	526
	Saving dashboards and charts.....	526
	Utilization Threshold charts.....	527
	Charts View.....	527
	Customizing a chart.....	528
	Customizing the tabbed Charts view.....	530
	Editing charts.....	532
	Copying a chart.....	533
	Analyze view.....	533
	Creating a dashboard from an Analyze view.....	534
	Creating a template dashboard from an Analyze view.....	535
	Changing the time range.....	535
	Symmetrix systems view (Real Time).....	536
	FE Director view (Real Time).....	537
	BE Director (DA) view (Real Time).....	537
	External Director view (Real Time).....	538
	RDF Director view (Real Time).....	538
	Array systems view (Diagnostic) .....	539
	Alerts view (Diagnostic).....	539
	FE Directors view (Diagnostic).....	540
	BE Directors (DA) view (Diagnostic).....	541
	External Directors view (Diagnostic).....	541
	RDF Directors view (Diagnostic).....	542
	IM Directors view (Diagnostic).....	542
	EDS Directors view (Diagnostic).....	543
	Cache Partitions view (Diagnostic).....	543
	Boards view (Diagnostic).....	543
	Disk Technologies view (Diagnostic).....	544
	Events view (Diagnostic).....	544
	Storage Groups view (Diagnostic).....	545
	Device Groups view (Diagnostic).....	546
	Databases view (Diagnostic).....	547
	Thin Pools view (Diagnostic).....	547

Disk Groups view (Diagnostic).....	548
External Disk Groups view (Diagnostic).....	549
SRPs view (Diagnostic).....	549
RDFA Groups view (Diagnostic).....	550
RDFS Groups view (Diagnostic).....	550
Snap Pools view (Diagnostic).....	551
DSE Pools view (Diagnostic).....	551
FAST VP Policies view (Diagnostic).....	552
Disk Group Tiers view (Diagnostic).....	552
Virtual Pool Tiers view (Diagnostic).....	553
Storage Groups view (Diagnostic).....	554
Hosts view (Diagnostic).....	555
Initiators view (Diagnostic).....	555
Masking Views view (Diagnostic).....	556
Port Groups view (Diagnostic).....	556
Host IO Limit by SG view (Diagnostic).....	556
Host IO Limit by FE view (Diagnostic).....	557
Symmetrix systems view (Historical).....	557
Alerts view (Historical).....	558
FE Directors view (Historical).....	559
BE Directors (DA) view (Historical).....	559
External Directors view (Historical).....	560
RDF Directors view (Historical).....	560
IM Directors view (Historical).....	561
EDS Directors view (Historical).....	561
Cache Partitions view (Historical).....	562
Boards view (Historical).....	562
Disk Technologies view (Historical).....	562
Events view (Historical).....	563
Storage Groups view (Historical).....	564
Device Groups view (Historical).....	565
Databases view (Historical).....	565
Thin Pools view (Historical).....	566
Disk Groups view (Historical).....	567
External Disk Groups view (Historical).....	567
SRPs view (Historical).....	568
RDFA Groups view (Historical).....	568
RDFS Groups view (Historical).....	569
Snap Pools view (Historical).....	570
DSE Pools view (Historical).....	570
FAST VP Policies view (Historical).....	571
Disk Group Tiers view (Historical).....	571
Virtual Pool Tiers view (Historical).....	572
Storage Groups view (Historical).....	572
Hosts view (Historical).....	573
Masking Views view (Historical).....	573
Port Groups view (Historical).....	574
Host IO Limit by SG view (Historical).....	574
Host IO Limit by FE view (Historical).....	575
Initiators view (Historical).....	575
Heatmap.....	575
Viewing Heatmap Metrics Charts.....	576
Navigating from Heatmap to Analyze or Charts.....	577
Filtering heatmaps.....	577
Reports.....	577
Report operations.....	578

Creating performance reports.....	579
Copying performance reports.....	580
Creating queries using the Create Query wizard.....	580
Editing queries using the Edit Query wizard.....	582
Modifying performance reports.....	583
Deleting performance reports.....	584
Running performance reports.....	584
Scheduling performance reports.....	584
Cancelling a scheduled report.....	585
Copying performance reports.....	586
Viewing Real Time traces.....	586
Creating a Real Time trace.....	587
Modifying a Real Time Trace.....	588
Deleting a Real Time trace.....	588
Plan View.....	588
SRP projection dashboard.....	589
Thin pool projection dashboard.....	590
Viewing system registrations.....	591
Viewing system registration details.....	592
Registering storage systems.....	592
Removing a system registration.....	593
Viewing registered storage systems information.....	593
Changing registration details.....	593
Managing dashboard catalog.....	594
Viewing Performance databases.....	594
Viewing database details.....	595
Restoring a database.....	597
Backing up a database.....	597
Canceling a scheduled database backup.....	598
Editing a scheduled Performance database backup.....	598
Editing database retention settings.....	599
Deleting databases.....	599
Removing database backup files.....	600
Viewing Performance thresholds and alerts.....	600
Performance Threshold Alert operations.....	601
Creating a performance threshold alert.....	601
Editing a performance threshold alert.....	602
Deleting performance thresholds and alerts.....	602
Configuring SNMP notifications.....	602
Managing dashboard catalog.....	603
Configuring email notifications.....	603
About exporting and importing performance settings.....	604
Importing Performance settings.....	604
Exporting Performance settings.....	605
Exporting Performance Viewer settings.....	605
Metrics.....	606
BE Director (DA) metrics.....	611
BE Emulation metrics.....	613
BE Port metrics.....	614
Board metrics.....	614
Cache Partition metrics.....	615
DATA Volume metrics.....	616
Database metrics.....	618
Database by Pool metrics.....	624
Device Group metrics.....	625
Disk metrics.....	630

Disk Bucket metrics.....	631
Disk Group metrics.....	632
Disk Group tier metrics.....	633
Disk Technology metrics.....	634
DSE Pool metrics.....	635
DX Emulation metrics.....	636
DX Port metrics.....	636
EDS Director metrics.....	637
EDS Emulation metrics.....	637
External Director metrics.....	638
External Disk metrics.....	640
External Disk Group metrics.....	641
FAST VP Policy metrics.....	642
FE Director metrics.....	643
FE Emulation metrics.....	650
FE Port metrics.....	650
FE Port - FE metrics.....	651
FE Port - SG metrics.....	651
FICON Emulation metrics.....	651
FICON Emulation Threads metrics.....	652
FICON Port Threads metrics.....	652
Host metrics.....	652
Host IO Limit by FE metrics.....	653
Host IO Limit by SG metrics.....	653
IM Director metrics.....	653
IM Emulation metrics.....	654
Initiator metrics.....	654
Initiators by Port metrics.....	654
IP Interface metrics.....	655
iSCSI Target metrics.....	656
Masking View metrics.....	656
Metas metrics.....	657
Other - Pool Bound Volume metrics.....	664
Pool by Storage Group metrics.....	669
Port Group metrics.....	670
RDF Director metrics.....	671
RDF Emulation metrics.....	673
RDF Port metrics.....	673
RDF/A Group metrics.....	673
RDF/S Group metrics.....	677
SAVE Volume metrics.....	682
Snap Pool metrics.....	688
Spare Disk metrics.....	688
SRP metrics.....	690
Storage Group metrics.....	692
Storage Group by Pool metrics.....	699
Storage Group by Tier metrics.....	700
Thin Pool metrics.....	701
Thin Volume metrics.....	703
Tier by Storage Group metrics.....	709
Virtual Pool Tier metrics.....	710
Volume metrics.....	712
Viewing and managing metrics.....	718
Editing metrics.....	719
Metrics.....	719
Setting the time range for viewing data.....	831



	Creating a template dashboard from an Analyze view.....	831
	Filtering performance data.....	832
	Filtering object lists.....	832
	Filtering heatmaps.....	832
<b>Chapter 8</b>	<b>Database Storage Analyzer</b>	<b>833</b>
	Introduction.....	834
	Database collection and retention policy.....	834
	Mapping files.....	835
	Viewing Databases page.....	835
	Viewing Database Administration page.....	837
	Registering a monitored environment.....	839
	Adding an Oracle database.....	839
	Registering Monitored Environment - Advanced Options.....	840
	Adding monitored MS SQL server instances.....	841
	Editing monitored Oracle databases.....	843
	Editing monitored MS SQL server instances.....	844
	Starting statistics collection.....	845
	Stopping statistics collection.....	845
	Running device mapping.....	845
	Schedule device mapping.....	845
	Removing monitored environment instance.....	846
	Viewing the Performance Page.....	846
	Viewing the Analytics Page.....	849
	Viewing analytics details.....	853
	Viewing database storage details.....	853
	Viewing database details.....	854
	Adding hints.....	854
	Viewing hints.....	855
	Editing hints.....	856
	Enabling hints.....	857
	Disabling hints.....	857
	Removing hints.....	858
	Viewing hint logs.....	858
	Hint operations.....	859
<b>Chapter 9</b>	<b>VMware</b>	<b>861</b>
	Understanding Unisphere support for VMware.....	862
	Viewing vCenters and ESXi information.....	862
	Registering vCenter/ESXi.....	863
	Editing vCenter/ESXi.....	864
	Unregistering vCenter/ESXi.....	864
	Rediscover vCenter/ESXi.....	865
	Viewing ESXi server details.....	865
	Viewing ESXi server masking views.....	866
	Viewing ESXi server performance details.....	866
	Viewing ESXi server virtual machines details.....	867
	Viewing ESXi server virtual machine disks.....	868
<b>Chapter 10</b>	<b>System Management</b>	<b>871</b>
	Viewing Storage System details.....	872
	Setting system attributes.....	874
	Using the Emulation Management wizard.....	877
	Setting CPU I/O resource distribution.....	877

Setting logging level preferences .....	878
Understanding eNAS.....	878
Discovering eNAS control stations.....	879
Managing File storage.....	879
Provisioning storage for file .....	880
Launching Unisphere for VNX.....	882
Managing file storage groups.....	882
Managing file masking views.....	883
Viewing file systems.....	883
Viewing file system details.....	884
Viewing file system storage pools.....	885
Viewing file system storage pool details.....	885
Manage file storage alerts.....	887
Viewing the system audit log.....	889
Viewing Symmetrix audit log details.....	890
Viewing system hardware.....	891
Viewing available ports.....	891
Viewing back-end directors.....	892
Viewing back-end director details.....	893
Viewing external directors.....	893
Viewing external director details.....	894
Viewing system front-end directors.....	895
Viewing system front end director details.....	897
Viewing RDF directors.....	899
Viewing RDF director details.....	901
Viewing RDF director SRDF groups.....	903
Viewing IM directors.....	905
Viewing EDS directors.....	906
Viewing failed drives.....	907
Viewing mapped front-end volumes.....	907
Engenuity Warning dialog box.....	908
Converting directors.....	909
Associating directors with ports.....	910
Setting director port attributes.....	910
Associating directors and ports.....	912
Disassociating directors and ports.....	913
Enabling and disabling director ports.....	913
Performing system health checks.....	913
Naming storage systems.....	914
Replacing failed drives.....	914
Managing jobs.....	918
Making configuration changes safely.....	918
Understanding task persistence.....	919
Previewing jobs.....	920
Scheduling jobs.....	920
Running jobs.....	921
Rescheduling jobs.....	921
Modifying jobs.....	922
Reordering tasks within a job.....	922
Grouping jobs.....	922
Un-grouping jobs.....	922
Stopping jobs.....	923
Deleting jobs.....	923
Viewing the job list.....	923
Viewing job details.....	924
Understanding licenses.....	925

Installing licenses.....	926
Removing host-based licenses.....	927
Viewing Symmetrix entitlements.....	927
Viewing host-based licenses.....	928
Viewing license usage.....	928
Viewing license file.....	929
Viewing license file details.....	930
Understanding access controls.....	931
Opening access controls.....	931
Creating access groups.....	932
Adding access ID to access groups.....	932
Removing access IDs from access groups.....	933
Deleting access groups.....	933
Viewing access groups.....	934
Viewing access group details.....	934
Creating access pools.....	935
Modifying access pools.....	935
Deleting access pools.....	936
Viewing access pools.....	936
Viewing access pool details.....	937
Creating access control entries.....	937
Deleting access control entries.....	938
Viewing access control entries.....	938
Viewing access control entry details.....	939
Viewing access IDs.....	939
Viewing access pool volumes.....	940
Viewing access types.....	940
Access types.....	941
Modifying access types.....	942
Understanding dynamic cache partitioning.....	943
Enabling/Disabling dynamic cache partitioning.....	943
Creating dynamic cache partitions.....	943
Modifying dynamic cache partitions.....	944
Assigning dynamic cache partitions.....	945
Deleting dynamic cache partitions.....	945
Running in analyze mode.....	946
Viewing dynamic cache partitions.....	947
Viewing dynamic cache partition details.....	948
Viewing volumes assigned to dynamic cache partitions.....	950
System management - iSCSI.....	951
Creating an iSCSI target.....	952
Modifying an iSCSI target.....	953
Creating an IP interface.....	954
Editing an IP interface.....	954
Adding an IP route.....	955
Deleting an iSCSI target.....	955
Deleting an IP interface.....	955
Removing an IP route.....	956
Attaching an IP interface to an iSCSI target.....	956
Attaching an iSCSI target to an IP interface.....	957
Detaching an IP interface from an iSCSI target.....	958
Disabling an iSCSI target.....	958
Enabling an iSCSI target.....	959
Setting port flags.....	959
Viewing the iSCSI directors list.....	960
Viewing the iSCSI director details.....	961

## CONTENTS

Viewing IP interfaces list .....	961
Viewing IP interfaces details.....	962
Viewing iSCSI targets list.....	964
Viewing iSCSI target details.....	965
Viewing IP routes list.....	966
Viewing the IP routes details.....	966
Viewing iSCSI ports list.....	967
Viewing iSCSI ports details.....	968

# TABLES

1	Service level compliance rules.....	63
2	User roles and associated permissions.....	81
3	Permissions for Local Replication, Remote Replication and Device Management roles .....	82
4	Host/Host group flags.....	299
5	TimeFinder/Clone session options.....	366
6	TimeFinder/Snap session options.....	377
7	TimeFinder/Mirror session options.....	383
8	Array metrics.....	606
9	BE director (DA) metrics.....	611
10	BE emulation metrics.....	614
11	BE port metrics.....	614
12	Board metrics.....	614
13	Cache partition metrics.....	615
14	DATA volume metrics.....	617
15	Database metrics.....	618
16	Database by pool metrics.....	624
17	Device group metrics.....	625
18	Disk metrics.....	630
19	Disk bucket metrics.....	631
20	Disk group metrics.....	632
21	Disk group tier metrics.....	633
22	Disk technology metrics.....	634
23	DSE pool metrics.....	636
24	DX emulation metrics.....	636
25	DX port metrics.....	637
26	EDS director metrics.....	637
27	EDS director metrics.....	637
28	External director metrics.....	638
29	External disk metrics.....	640
30	External disk group metrics.....	641
31	FAST VP policy metrics.....	642
32	FE director metrics.....	643
33	FE director metrics.....	646
34	FE emulation metrics.....	650
35	FE port metrics.....	650
36	FE port (FE) metrics.....	651
37	FE port (SG) metrics.....	651
38	FICON emulation metrics.....	651
39	FICON emulation threads metrics.....	652
40	FICON port threads metrics.....	652
41	Host metrics.....	652
42	Host IO limit (by FE) metrics.....	653
43	Host IO limit (by SG) metrics.....	653
44	IM director metrics.....	653
45	IM emulations metrics.....	654
46	Initiator metrics.....	654
47	Initiators (by port) metrics.....	655
48	IP interface metrics.....	655
49	iSCSI target metrics.....	656
50	Masking view metrics.....	656
51	Metas metrics.....	657
52	Pool-bound volumes metrics.....	664

53	Pool by storage group metrics.....	669
54	Port group metrics.....	670
55	RDF director metrics.....	671
56	RDF emulation metrics.....	673
57	RDF port metrics.....	673
58	RDF/A group metrics.....	673
59	RDF/S group metrics.....	677
60	SAVE volume metrics.....	682
61	Snap pool metrics.....	688
62	Spare disk metrics.....	688
63	SRP metrics.....	690
64	Storage group metrics.....	693
65	Storage group (by pool) metrics.....	700
66	Storage group (by tier) metrics.....	700
67	Thin pool metrics.....	701
68	Thin volume metrics.....	704
69	Tier (by storage group) metrics.....	709
70	Virtual pool tier metrics.....	710
71	Volume metrics.....	712
72	Array metrics.....	719
73	BE director (DA) metrics.....	725
74	BE emulation metrics.....	727
75	BE port metrics.....	727
76	Board metrics.....	728
77	Cache partition metrics.....	728
78	DATA volume metrics.....	730
79	Database metrics.....	731
80	Database by pool metrics.....	738
81	Device group metrics.....	739
82	Disk metrics.....	744
83	Disk bucket metrics.....	745
84	Disk group metrics.....	745
85	Disk group tier metrics.....	746
86	Disk technology metrics.....	748
87	DSE pool metrics.....	749
88	DX emulation metrics.....	750
89	DX port metrics.....	750
90	EDS director metrics.....	750
91	EDS director metrics.....	751
92	External director metrics.....	751
93	External disk metrics.....	753
94	External disk group metrics.....	755
95	FAST VP policy metrics.....	756
96	FE director metrics.....	756
97	FE director metrics.....	759
98	FE emulation metrics.....	763
99	FE port metrics.....	763
100	FE port (FE) metrics.....	764
101	FE port (SG) metrics.....	764
102	FICON emulation metrics.....	765
103	FICON emulation threads metrics.....	765
104	FICON port threads metrics.....	765
105	Host metrics.....	765
106	Host IO limit (by FE) metrics.....	766
107	Host IO limit (by SG) metrics.....	766
108	IM director metrics.....	767

109	IM emulations metrics.....	767
110	Initiator metrics.....	767
111	Initiators (by port) metrics.....	768
112	IP interface metrics.....	768
113	iSCSI target metrics.....	769
114	Masking view metrics.....	769
115	Metas metrics.....	770
116	Pool-bound volumes metrics.....	777
117	Pool by storage group metrics.....	782
118	Port group metrics.....	783
119	RDF director metrics.....	784
120	RDF emulation metrics.....	786
121	RDF port metrics.....	786
122	RDF/A group metrics.....	786
123	RDF/S group metrics.....	790
124	SAVE volume metrics.....	795
125	Snap pool metrics.....	801
126	Spare disk metrics.....	801
127	SRP metrics.....	803
128	Storage group metrics.....	805
129	Storage group (by pool) metrics.....	812
130	Storage group (by tier) metrics.....	813
131	Thin pool metrics.....	814
132	Thin volume metrics.....	816
133	Tier (by storage group) metrics.....	822
134	Virtual pool tier metrics.....	823
135	Volume metrics.....	825
136	Task status before and after server shutdown.....	919
137	Access types.....	941





# CHAPTER 1

## Introduction

- [Unisphere Online Help](#).....26
- [Capacity information](#).....29

# Unisphere Online Help

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

A HTML5 based Unisphere provides a number of advantages:


- improved security
- modern user interface
- reduced application response times

Unisphere supports the following tasks which are available from the items on the side panel and blue title bar:


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

- **HOME**- View system view dashboard of all storage systems being managed.
- **PERFORMANCE** - Monitors and manages storage system performance data (Dashboards, Charts, Analyze, Heatmap, Reports, Plan, Real Time traces, and Performance Database management). Refer to [Understanding Performance Management](#) on page 518 for more information.
- **VMWARE**- Views all the relevant storage related objects to an ESXi server and also provides the ability to help troubleshooting storage performance related issues to the ESXi server. Refer to [Understanding Unisphere support for VMware](#) on page 862 for more information.
- **DATABASES**- Monitors and troubleshoots database performance issues. Refer to [Introduction](#) on page 834 for more information.
- **EVENTS**- Includes Alerts and Job List.
- **SUPPORT**- Displays support information.
- 



Click  to set preferences. Refer to for more information.



Click  to hide the side panel and click it again to display the side panel.

Click **HOME** to return to the **All Systems** view.

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

- **HOME**- View system view dashboard of all storage systems being managed.
- **DASHBOARD**- View the following dashboards for a selected storage system: Capacity and Performance, System Health, Storage Group compliance, Capacity, and Replication.
- **STORAGE**- Manage storage (storage groups, service levels, templates, SRPs, volumes, external storage, VVols, FAST policies, tiers, thin pools, disk groups and VLUN migration). Refer to [Understanding Storage Management](#) on page 90 for more information.


- **HOSTS**- Manage hosts (hosts, masking views, port groups, initiators, XtremSW Cache Adapters, PowerPath Hosts, Mainframe , and CU images). Refer to [Understanding Host Management](#) on page 292 for more information.
- **DATA PROTECTION** -Manage data protection (storage groups, device groups, SRDF groups, migrations, virtual witness, open replicator, SRDF/A DSE pools, TimeFinder SNAP pools, and RecoverPoint systems). Refer to [Understanding Data Protection Management](#) on page 348 for more information.
- **PERFORMANCE** - Monitors and manages storage system performance data (Dashboards, Charts, Analyze, Heatmap, Reports, Plan, Real Time traces, and Performance Database management). Refer to [Understanding Performance Management](#) on page 518 for more information.
- **SYSTEM**- Includes Hardware, Symmetrix Properties, File (eNAS), and iSCS..
- **EVENTS**- Includes Alerts, Job List, and Audit log.
- **SUPPORT**- Displays support information.

### New and modified features/functionality in 9.0.0

- **HTML5 support** -A HTML5 based Unisphere provides a number of advantages:
  - improved security
  - modern user interface look and feel - use of browser functionality, bookmarks for links, back and forward buttons. Facilitates enhanced collaboration as you can share links to system views with colleagues.
  - reduced application response times
  - aligns with other Dell EMC products
- **System Health Score** -The System Health dashboard provides a single place from which you can quickly determine the health of the system. The System Health panel displays values for the following high level health or performance metrics: Configuration, Capacity, System Utilization, Storage Group Response Time and Service Level Compliance. It also displays an overall health score based on the lowest health score out of the five metrics. These five categories are for storage systems running HYPERMAX OS 5977 or higher. For storage systems running Enginuity OS 5876, the health score is based on four categories: Configuration, System Utilization, Capacity and storage group (SG) Response Time. The health score is calculated every five minutes. The overall value is always calculated from all metric values. If a health score category is seen as stale or unknown then the overall health score is not updated. The previously calculated overall health score is displayed but its value is denoted as stale by setting the menu item to grey (refer to [Understanding the system health score](#) on page 37).
- **Role Based Access Control (RBAC)** - This feature provides a set of roles with more granular access that can be assigned to users in order to limit what resources can be accessed and what functions a user can perform on those resources. Additional roles are Device Management, Local Replication and Remote Replication at the entire array or an storage group subset. This feature also supports, for tracking purposes, a full audit log of users and actions performed. (refer to [Adding authorization rules](#) on page 72).
- **Service Levels (Performance QoS)** - Unisphere supports all service levels (Diamond, Platinum, Gold, Silver, Bronze and Optimized) for FBA SRPs containing internal disk groups on Storage systems running PowerMaxOS 5978 and above. There are no changes to service level restrictions for CKD SRPs or SRPs containing external disk groups (refer to [Viewing service levels](#) on page 152).
- **Compliance** - A Compliance Tab has been added to the the Storage Group detailed view page (refer to [Viewing Storage Group Compliance view](#) on page 94).

- Noisy Neighbors - The Noisy Neighbors feature displays the following performance data for a selected storage group:
  - FE Directors details - Name, % busy, and queue depth utilization.
  - FE Port details - Name, % busy, and host I/Os per second.
  - Related SGs - Name, response time, host I/Os per second, and host MBs per second.
 (refer to [Viewing ESXi server performance details](#) on page 866).
- Data Reduction - Data is reduced using data compression and de-duplication (de-duplication applies for storage systems running PowerMaxOS 5978 or higher).
- Real Time Data Collection - This feature provides the ability to troubleshoot at a more granular level for a set number of Storage Group's for a limited set of metrics at a 30 second level. This will be limited to 1 array per time, a maximum of 5 SG's at a time and a certain number of KPI metrics. The metrics reported on are Response Time, Host I/O's Per Sec, Host MB's Per Sec, Host Reads Per Sec, and Host Writes Per Sec.
- SRDF and Metro topology view - The SRDF and Metro topology view visually describes the layout of the SRDF connectivity of the selected storage system in Unisphere.
- Storage Templates - Using the configuration and performance characteristics of an existing storage group as a starting point, you can create templates that will pre-populate fields in the provisioning wizard and create a more realistic performance reservation in your future provisioning requests (refer to [Creating storage templates](#) on page 267).
- VMware integration - Unisphere support for VMware provides the storage admin access to all the relevant storage related objects to an ESXi server and also provides the ability to help troubleshooting storage performance related issues to the ESXi server. You can, as a read only user, discover at the vCenter level as well as discovering an individual ESXi server. If a vCenter is discovered, then all ESXi servers under that vCenter are discovered. All ESXi servers, that do not have local storage on the Unisphere performing the discovery, are filtered out. Once VMware information is added by a user, all other users of Unisphere are able to access this information. The minimum version number supported by vCenter is version 5.5. The VMware feature supports a maximum of 75 ESXi servers and 2000 VMs per Unisphere for PowerMax install (refer to [Understanding Unisphere support for VMware](#) on page 862).
- Integration of Database Storage Analyzer into Unisphere - DSA for Oracle and SQL now fully integrated with Unisphere, no separate login or page launch required. The DB mapping procedure has been streamlined to make it more user friendly (refer to [Introduction](#) on page 834).
- Silent Install - This supports installations of Unisphere by invoking a automated script which handles the various steps involved. Included is a response file containing default values that the user can edit. Where there is not enough space or memory on a host, the install will be aborted.

### Using the help system

Clicking  on the navigation bar results in the display of three options. Clicking the top option results in the display of a window that displays the Unisphere help home page. Clicking the middle option results in the display of a window that displays the Unisphere help for that screen (context-sensitive help). Clicking the bottom option (**About**) results in the display of a window that displays the Unisphere version number.

Finding information:

- Using the Contents tab—Click the book icon to expand the table of contents and display help topics.
- Using the Search tab—Click the Search tab in the navigation pane. Type a search word or phrase and a list of topics that contain the word or phrase displays in the navigation panel. Click on the name of the topic to display it in the View panel.

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

### Supporting documentation

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

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

## Capacity information

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

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

Name	Abbreviation	Binary Power	Binary Value (in Decimal)	Decimal Power	Decimal (Equivalent)
kilobyte	KB	2 <sup>10</sup>	1,024	10 <sup>3</sup>	1,000
megabyte	MB	2 <sup>20</sup>	1,048,576	10 <sup>6</sup>	1,000,000
gigabyte	GB	2 <sup>30</sup>	1,073,741,824	10 <sup>9</sup>	1,000,000,000
terabyte	TB	2 <sup>40</sup>	1,099,511,627,776	10 <sup>12</sup>	1,000,000,000,000



# CHAPTER 2

## Getting Started

- [Operating as the initial setup user.....](#) 32
- [Viewing Home Dashboard view - All Storage Systems.....](#) 32
- [Viewing Home Dashboard view - Specific Storage System.....](#) 34
- [Viewing system performance view.....](#) 34
- [Viewing the System Health Dashboard.....](#) 35
- [Understanding the system health score.....](#) 37
- [Viewing Storage Group Compliance view.....](#) 39
- [Viewing Capacity dashboard view.....](#) 40
- [Viewing Replication dashboard.....](#) 41
- [Discovering storage systems.....](#) 43
- [Refreshing storage system information.....](#) 44
- [Viewing product version information.....](#) 44
- [Searching for storage objects.....](#) 45
- [Modifying server logging levels.....](#) 46
- [Exiting the console.....](#) 46
- [Getting help.....](#) 46

## Operating as the initial setup user

When Unisphere 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 SecurityAdmin 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.

When logging in to Unisphere as the Initial Setup User (ISU), the "Initial setup user warning" message is displayed and it informs you that you can only access the listed storage systems because they do not have defined authorization rules. Once rules are defined for the storage systems, you will no longer be able to access or view the storage systems as the ISU.

To continue to access/view a storage system while operating as the ISU, select the corresponding Assign Admin role to ISU option and click OK.

The main tasks of an ISU are:

- [Creating local users](#) on page 76
- [Adding authorization rules](#) on page 72

For more information on operating as the ISU, refer to the *Unisphere for PowerMax Installation Guide*.

## Viewing Home Dashboard view - All Storage Systems

### Before you begin

The user requires a minimum of Monitor permissions to perform this task.

For an overview of Unisphere functionality, see [Unisphere Online Help](#) on page 26.

The Home Dashboard view (the default mode on login) provides an overall view of the status of all of the storage systems managed by Unisphere. The following panels are displayed:

- Compliance
- Capacity
- Health score - see [Understanding the system health score](#) on page 37
- Throughput
- IOPS
- Efficiency

You can sort the view by the following:

- Compliance
- Capacity
- Health score - see [Understanding the system health score](#) on page 37
- Throughput
- IOPS
- Efficiency

To view the home dashboard view:



## Procedure

1. From the main menu, click **Unisphere for PowerMax**.


The home dashboard view is displayed.

2. View the following parameters displayed in each storage system panel.

Depending on the metric selected, some of the following items are displayed:



- Storage system ID - The serial number of the storage system.
- Storage system model - The model number of the storage system.
- The version of microcode on the storage system.
- Data chart - The information displayed in the chart depends on the selected metric.
- Capacity - Percentage of currently allocated capacity for the storage system.
- Compliance - Service level compliance data in the form of Storage group counts for each compliance state (Critical, Marginal, Stable) as well as total Storage Group count and number of Storage groups with no service level assigned.
- Performance - Current performance health score.
- Throughput - Current throughput for the system, in MB/second.
- IOPS - Current IOPS for the system.
- Efficiency - The overall efficiency ratio for the array. It represents the ratio of the sum of all TDEVs plus snapshot sizes (calculated based on the 128K track size) and the physical Used Storage (calculated based on the compressed pool track size).

3. (Optional) To view the alerts, click  on any storage system panel and click **VIEW ALERTS**.

The  color reflects the highest severity alert for the associated storage system.

You can also view the job list and navigate to the compliance view for the storage system by clicking the related icon on any storage system panel.

Note: you can view alerts for remote storage systems and storage systems that are not registered to collect performance data.

4. (Optional) Click  to view the system view in list format.
5. (Optional) Click  to view the system view in card view format.
6. (Optional) From a panel view or list view, click the storage system identity of the system you want to view in more detail.
7. (Optional) To navigate to other areas, click on any of the following from the left hand panel:
  - **HOME**
  - **PERFORMANCE**
  - **VMWARE**

- DATABASES
- EVENTS
- SUPPORT

## Viewing Home Dashboard view - Specific Storage System

### Before you begin

The user requires a minimum of Monitor permissions to perform this task.

For an overview of Unisphere functionality, see [Unisphere Online Help](#) on page 26.

The Home Dashboard view for a specific storage system provides a view of the status of a specific storage system managed by Unisphere. The following panels are displayed:

- PERFORMANCE
- SYSTEM HEALTH
- SG COMPLIANCE
- CAPACITY
- REPLICATION

To view the home dashboard view:

### Procedure

1. From the main menu, click **Unisphere for PowerMax**.

The home dashboard view for all storage systems is displayed.

2. Select a storage system.

The system performance dashboard is displayed by default (see [Viewing system performance view](#) on page 34).

3. (Optional) To navigate to other areas, click on any of the following from the left hand panel:

- HOME
- DASHBOARD
- STORAGE
- HOSTS
- DATA PROTECTION
- PERFORMANCE
- SYSTEM
- EVENTS
- SUPPORT

## Viewing system performance view

### Before you begin

The user requires a minimum of Monitor permissions to perform this task.

To view the system performance view:

### Procedure

1. From the main menu, click **Unisphere for PowerMax**.
2. Select a storage system.  
The home dashboard view for the selected storage system is displayed. The system performance dashboard is displayed by default.
3. View the capacity and performance data for the selected storage system. The following items are displayed:
  - A **Capacity** panel displaying the following:
    - A graphical representation of the system's subscribed and usable capacity (used = blue and free = grey) and the percentage used for both.
    - The percentage of subscribed usable capacity.
    - The overall efficiency ratio.
  - A **Performance** panel displaying the following graphs over a four hour, one week, or two weeks period:
    - **Host IOs per sec** in terms of read and write operations over time
    - **Latency** in terms of read and write operations over time
    - **Throughput** in terms of read and write operations over time
  - A **Capacity Trend** panel displaying usable capacity and subscribed capacity in terabytes.
  - The following control is available:
    - **VIEW PERFORMANCE DASHBOARD**—[Viewing dashboards](#) on page 519
4. (Optional) To navigate to other areas, click on any of the following from the left hand panel:
  - HOME
  - DASHBOARD
  - STORAGE
  - HOSTS
  - DATA PROTECTION
  - PERFORMANCE
  - SYSTEM
  - EVENTS
  - SUPPORT

## Viewing the System Health Dashboard

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

The System Health section displays values for the following five high level health or performance metrics: System Utilization, Configuration, Capacity, SG Response Time and Service Level Compliance. It also displays an overall health score based on these

five categories. The overall system health score is based on the lowest health score out of the categories System Utilization, Configuration, Capacity, SG Response Time and service level compliance. See [Understanding the system health score](#) on page 37 for details on how these scores are calculated. These five categories are for systems running HYPERMAX OS 5977 or later. For systems running Enginuity 5876, the health score is based on the Hardware, Configuration, Capacity and SG Response time scores. The health score is calculated every five minutes. Note: The Health score values for Hardware, SG Response and service level compliance are not real-time; they are based on values within the last hour.

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

To view the system health dashboard:

#### Procedure

1. Select the storage system.
2. Optional: Hover over **SYSTEM HEALTH** to view the system health summary for the storage system.
3. Click **SYSTEM HEALTH** and view the following items:
  - **Introducing your Health Score** - [Understanding the system health score](#) on page 37
  - **Health Score** panel - The current score and the 30 day trend are displayed for the storage system health parameters - Total Issues, Configuration, Capacity, System Utilization, Service Level Compliance, and SG Response Time.  
The following views are available by clicking on the associated panel item:
    - **VIEW ALERTS**—[Viewing alerts](#) on page 52
    - **VIEW PERFORMANCE**—[Using default dashboards](#) on page 520
    - **VIEW STORAGE GROUPS**— [Viewing storage groups](#) on page 135
  - **Hardware** panel - The storage system hardware is displayed in terms of the number of front end (FE) directors, SRDF directors, back end (BE) directors, available ports and cache partitions.  
The following views are available by clicking on the associated panel item:
    - **Front End**—[Viewing system front-end directors](#) on page 895
    - **RDF**—[Viewing RDF directors](#) on page 899
    - **Back End**—[Viewing back-end directors](#) on page 892
    - **Available Ports** —[Viewing available ports](#) on page 891
  - The following controls are available:
    - **VIEW SYMMETRIX PROPERTIES**—[Viewing Storage System details](#) on page 872
    - **MANAGE EMULATION**—[Using the Emulation Management wizard](#) on page 877 (For storage systems running HYPERMAX OS 5977 or higher)
    - **VIEW RESERVATIONS**—[Viewing reservations](#) on page 277
    - **VIEW OTHER HARDWARE**—[Viewing dynamic cache partitions](#) on page 947

- **RUN HEALTH CHECK**—[Performing system health checks](#) on page 913
- **RUN DISK REPLACEMENT**—[Replacing failed drives](#) on page 914 (For storage systems running Enginuity OS 5876)

## Understanding the system health score

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

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

The Configuration health score is based on storage system hardware alerts in the system like Director and Port alerts. The System Utilization, Capacity, storage group response time and service level compliance are based on performance information.

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

The Configuration health score is calculated as follows:

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

The Capacity health Score is based on percentage of used usable. Capacity levels are checked at the Array, SRP (only on storage systems running HYPERMAX OS 5977 or higher) and Thin Pool level (only storage systems running Enginuity OS 5876).

The capacity health scores are calculated as follows:

- Critical level: 95% - 30 points reduced
- Warning level: 80% - 10 points reduced

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

- FE\_DIR: PERCENT\_BUSY, QUEUE\_DEPTH\_UTILIZATION
- FE\_PORT: PERCENT\_BUSY
- BE\_PORT\_DA: PERCENT\_BUSY
- BE\_DIR\_DA: PERCENT\_BUSY
- RDF\_PORT: PERCENT\_BUSY
- RDF\_DIR: PERCENT\_BUSY

- BE\_PORT\_DX: PERCENT\_BUSY
- BE\_DIR\_DX: PERCENT\_BUSY
- IM\_DIR: PERCENT\_BUSY
- EDS\_DIR: PERCENT\_BUSY
- BOARD: UTILIZATION
- CP: WP
- DISK: PERCENT\_BUSY

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

FE Port - Percent Busy - Critical 70, Warning 50

FE Director - Percent Busy - Critical 70, Warning 50; Queue Depth Utilization - Critical 75, Warning 60

BE Port DA - Percent Busy - Critical 70, Warning 55

BE Director DA - Percent Busy - Critical 70, Warning 55

RDF Port - Percent Busy - Critical 70, Warning 50

RDF Director - Percent Busy - Critical 70, Warning 50

BE Port DX - Percent Busy - Critical 70, Warning 55

BE Director DX - Percent Busy - Critical 70, Warning 55

IM Director - Percent Busy - Critical 70, Warning 55

EDS Director - Percent Busy - Critical 70, Warning 55

Board - Utilization - Critical 70, Warning 60

Cache Partition - Percent Busy - Critical 75, Warning 55

Disk - Percent Busy - Critical 70, Warning 55

The system utilization score is calculated as follows:

- Critical level: - 30 points reduced
- Warning level: - 10 points reduced

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

It is based on WLP Workload state. A reduction from the health score is performed when storage groups, which have a service level defined, are not meeting the service level requirements.

The Service Level compliance score is calculated as follows:

- Underperforming: - 30 points reduced
- Marginal performing: - 10 points reduced

Storage systems running Enginuity OS 5876: The Storage Group Response health score is based on software category health scores. Certain key metrics are examined against threshold values and if they exceed a certain threshold, then the health score is negatively affected

The storage group response score is calculated as follows:

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

- Read Response Time: Critical: 30 points reduced; Warning : 20 points reduced
- Write Response Time:
- Response Time: 30 points reduced; Warning : 20 points reduced
- Database: Read Response Time, Write Response Time, Response Time

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

## Viewing Storage Group Compliance view

### Before you begin

The user requires a minimum of Monitor permissions to perform this task.

To view the Storage Group (SG) Compliance view:





### Procedure

1. Select a storage system.

The system performance dashboard is displayed by default (see [Viewing system performance view](#) on page 34).

2. Optional: Hover over **SG COMPLIANCE** to view the storage health summary for the storage system.

3. Click **SG COMPLIANCE** and view the following items:

- **Compliance** panel—Displays how well the storage system's workload is complying with the overall service level. Storage groups compliance information displays for storage systems registered with the Performance component. The total number of storage groups is listed, along with information about the number of storage groups performing according to service level targets. Possible values are:
  -  **Critical**—Number of storage groups performing well below service level targets.
  -  **Marginal**—Number of storage groups performing below service level targets.
  -  **Stable**—Number of storage groups performing within the service level targets.
  -  **No Status** —Number of storage groups without a status.
- **Storage Groups** panel—The storage groups are listed in a view that can be filtered.
- The following controls are available:
  - **VIEW COMPLIANCE REPORT**—[Viewing compliance reports](#) on page 157(For storage systems running HYPERMAX OS 5977 or higher)
  - **VIEW ALL STORAGE GROUPS**—[Viewing storage groups](#) on page 135
  - **VIEW FAST STORAGE GROUPS**—[Viewing FAST storage groups](#) on page 173 (For storage systems running Engenuity OS 5876)

- **PROVISION STORAGE**—[Using the Provision Storage wizard](#) on page 100
- **EXCLUDE DATA**—[Managing Data Exclusion Windows](#) on page 158 (For storage systems running HYPERMAX OS 5977 or higher)

## Viewing Capacity dashboard view

### Before you begin

The user requires a minimum of Monitor permissions to perform this task.

To view the Capacity dashboard view:

### Procedure

1. Select a storage system.

The system performance dashboard is displayed by default (see [Viewing system performance view](#) on page 34).

2. Optional: Hover over **CAPACITY** to view the capacity summary for the storage system.
3. Click **CAPACITY**, select the system or an SRP instance (not applicable for systems running Enginuity 5876) and view the following items:

System running Enginuity 5876:

A graphical representation of the system's physical and virtual capacity (used = blue and free = grey) and the percentage used for both.

System running HYPERMAX OS 5977 or PowerMaxOS 5978 - System selected and **Show Detailed** selected :

- A graphical representation of the system's subscribed, snapshot and usable capacity (used = blue and free = grey) and the percentage used for both.
- A textual representation of the system's subscribed usable capacity.
- **System Usage** is displayed if you turn on the **Show Detailed** slider. The information is displayed in terms of System Meta data used, Replication Meta Data and Front End Meta Data. You can click on **Analyze Trend** to analyze trends across metrics particular to capacity and usage. Trending shown for Metadata usage, Subscribed Capacity, Snapshot Capacity, and Usable Capacity.
  - Metadata trending will capture System, Replication, Front-end and Back-end.
  - Subscribed capacity trending will capture all (non-shared and shared) allocated against total subscribed capacity.
  - Snapshot capacity trending will capture all (shared and non-shared) modified capacity against total snapshot capacity.
  - Usable capacity trending will capture all (user, system, temp) used capacity against total usable capacity.

---

### Note

The data shown depends on the code level the system is running. Front-end meta data isn't shown for systems running HYPERMAX OS 5977.

---



- **Efficiency** is also displayed in terms of Overall Efficiency Ratio, Data Reduction (Ratio and Enabled Percent), Virtual Provisioning savings, and Snapshot savings.

System running HYPERMAX OS 5977 or higher - SRP instance selected:

- A graphical representation of the system's subscribed, snapshot and usable capacity (used = blue and free = grey) and the percentage used for both.
- A textual representation of the system's subscribed usable capacity.
- Headroom details are also displayed. Headroom is displayed by default as an overall figure but can also be filtered to display headroom for OLTP, OLTP + Replication, DSS, DSS + Replication, and None.  
The headroom displayed depends on the system's code level — systems running HYPERMAX OS 5977 will only show Diamond service levels and a combination of workload types. Systems running PowerMaxOS 5978 will show headroom for the different SLO types (workload types aren't supported by this code level).
- **Efficiency** is also displayed in terms of Overall Efficiency Ratio, Data Reduction (Ratio and Enabled Percent), Virtual Provisioning savings, and Snapshot savings.

The following controls are available (for storage system running HYPERMAX OS 5977 or higher when a SRP instance is selected) from the **Actions** panel:

- **STORAGE GROUP DEMAND**—[Viewing Storage Group Demand Reports](#) on page 91
- **SERVICE LEVEL DEMAND**—[Viewing Service Level Demand Reports](#) on page 92
- **COMPRESSIBILITY**—[Viewing compressibility reports](#) on page 288

## Viewing Replication dashboard

### Before you begin

The user requires a minimum of Monitor permissions to perform this task.

The Replication Dashboard provides Storage Group Summary Protection information, summarizing worst states of various Replication technologies and counts of Management objects participating in these technologies. Storage systems running Enginuity 5876 also display a Device Group Summary, with counts of various Replication Technologies using Device Groups.

To view the Replication dashboard:

### Procedure

1. From the main menu, click **Unisphere for PowerMax**.
2. Select a storage system.

The system performance dashboard is displayed by default (see [Viewing system performance view](#) on page 34).

3. Click **REPLICATION** and view the following items:
  - A **Storage Group Summary** panel is displayed. For systems running HYPERMAX OS 5977 and higher, summary information for SRDF, SRDF/Metro and SnapVX is displayed. For systems running Enginuity 5876,

summary information for SRDF and Device Groups is displayed. To view the Storage Groups that are in the states indicated, you can click on the row which brings you to the technologies Storage Group list view that is filtered to show only the applicable Storage Groups for the selected state.

- A visual display of SRDF topology:

The SRDF Topology View visually describe the layout of the SRDF connectivity of the selected storage system in Unisphere.

It calculates this with a maximum of two hops, for example, Symm A has SRDF Groups to Symm B, which has SRDF Groups to Symm C, if a fourth storage system Symm D has SRDF Groups to Symm C but is not connected to Symm A or Symm B, it is not shown as it is outside the two hop count for the Array that Unisphere is currently managing. All types of SRDF Groups are used to calculate this view.

There are two components which make up the topology view, nodes and edges. A node is the storage system and the edges are the connectivity between the storage systems.

The edges are color coded in the familiar traffic light system. The colors are Green, Yellow and Red.

Green edges indicate that the state of the connectivity between two nodes (Arrays) is Good.

Yellow indicates that the connectivity between two nodes is degraded. Degraded in this case means that one or more SRDF Groups between the 2 arrays are either in a Transmit Idle state, or have some ports in an SRDF Group that are offline.

Red indicates that the state of the connectivity between two nodes is Critical. Critical in this case means that one or more SRDF Groups between the 2 arrays are Offline, or one more SRDF Groups contains ports that are all offline.

The edges are also drawn differently depending on the modes of the SRDF Groups between two arrays. A legend is available under the view:

Edges that are drawn with short dashes and short gapes between the dashes indicate that all the SRDF Groups between the two Arrays are Metro or Synchronous SRDF Groups.

Edges that are drawn with longer dashes and a short gap between the dashes indicate that all the SRDF Groups between the two Arrays are Asynchronous

Edges that are solid indicate that there is a mix of Asynchronous, Synchronous and SRDF/Metro SRDF Groups between the two Arrays.

Edges that are drawn with short dashes and a long distance between the dashes indicate the SRDF Groups between the two Arrays are Other SRDF Groups than mentioned above, including Empty SRDF Groups, Virtual Witness, Adaptive Copy, etc.

The nodes are drawn with some basic information about the Array, including Symmetrix ID and if set, the user defined nice name of the Array. An icon specific to the model of the Array will also be drawn into the node.

- A visual display of Migration Environments:

The Migrations Environments topology view visually describes the layout of the migration environments of the currently selected storage system.

The edges are color coded using the familiar traffic light system Red, Yellow and Green, Red in this case meaning the Migration Environment is in an Invalid State and Green meaning it is in a valid state. The color of the edge can also be dictated by the worst state of any migrations using this Environment. The edges are all drawn in a solid full line.

The nodes are drawn with some basic information about the storage system, including ID and if set, the user defined nice name of the storage system. An icon specific to the model of the Array is drawn into the node.

- Both topology views have the following controls:



Re-center - this brings the topology views nodes and edges back into full view.



- Zoom the view in to see nodes and edges.



- zoom the view out to view more of the topology view.



- allows the user view a fuller screen view of the topology. Clicking this opens a popup that takes up most of the visible space on the screen. All functionality of this view is the same as the view embedded on the page.



Clicking displays a Layout Manager which provides layout change options.

- The following controls are available:
  - **CREATE SNAPSHOT**—[Creating snapshots](#) on page 387 (For storage systems running HYPERMAX OS 5977 or higher)
  - **CREATE SRDF GROUP**—[Creating SRDF groups](#) on page 457
  - **CREATE MIGRATION ENVIRONMENT**—[Setting up a migration environment](#) on page 510
  - **CREATE VIRTUAL WITNESS**—[Adding SRDF Virtual Witness instances](#) on page 445 (For storage systems running HYPERMAX OS 5977 or higher)
  - **CREATE DEVICE GROUP**—[Creating device groups](#) on page 348 (For storage systems running Enginuity OS 5876)

## 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:

**Procedure**


1. On the **HOME** view, click **Unisphere for PowerMax** in the title bar.
2. Select a storage system
3. Click the arrow next to the storage system ID in the title bar and select **DISCOVER SYSTEMS**.
4. Read the warning stating the operation may take some time and click **OK** to confirm if you wish to proceed.

## Refreshing storage system information

Unisphere refreshes all of the storage system data from its database. This operation does not discover new storage systems, only refreshes data for existing systems.

To refresh a storage system:

**Procedure**

1. Select a storage system.
2. In the **Dashboard**, click  in the title bar.
3. Click **OK** to the **System Refresh Confirmation** dialog.

## Viewing product version information

**Procedure**

1. Select **SUPPORT** to open the **Support** view.

The following Latest Software properties display:

- Installed Unisphere Version
- Latest Available Unisphere Version
- Installed Solution Enabler Version
- Latest Available Solution Enabler Version

The following Solution Enabler properties display:

- **Connection Type**—Connection Type
- **Net Connection Security Level**—Net Connection Security Level
- **Net Protocol** —Net Protocol
- **Net Address** —Net Address
- **Net Port**—Net Port
- **Node Name**—Node Name
- **OS Type** —OS Type
- **OS Name** —OS Name
- **OS Version** —OS Version
- **OS Release** —OS Release
- **Machine Type** —Machine Type
- **System Time**—System Time

- **Num Symm Pdevs**—Number of Symm Pdevs
- **SYMAPI Build Version**—SYMAPI Build Version
- **SYMAPI Runtime Version** —SYMAPI Runtime Version
- **Library Type**—Library Type
- **64 bit Libraries**—64 bit Libraries
- **Multithread Libraries**—Multithread Libraries
- **Server Processor**—Server Processor
- **Storage Daemon**—Storage Daemon
- **GNS** —GNS
- **Storage Daemon GK Mgmt** —Storage Daemon GK Mgmt
- **Storage Daemon Caching**—Storage Daemon Caching
- **Storage Daemon Emulation**—Storage Daemon Emulation
- **Storage Daemon EM Caching** —Storage Daemon EM Caching
- **VMware Guest**—VMware Guest
- **Type of SYMAPI Database** —Type of SYMAPI Database
- **SYMAPI Lib Version which discovered DB** —SYMAPI Lib Version which discovered DB
- **SYMAPI Lib Version which wrote DB** —SYMAPI Lib Version which wrote DB
- **Minimum Edit Level of SYMAPI Lib Required**—Minimum Edit Level of SYMAPI Lib Required
- **Database Sync Time** —Database Sync Time
- **DG Modify Time**—DG Modify Time
- **Device in Multiple Device Groups**—Device in Multiple Device Groups


The following operations are available from the **Actions** panel:

- **PRODUCT SUPPORT PAGE**—clicking this brings you to the product support page.
- **SERVICE CENTER**—clicking this brings you to the service center.
- **MODIFY SERVER LOGGING** - [Modifying server logging levels](#) on page 46.

## Searching for storage objects

This procedure explains how to search for objects (storage groups, hosts, initiators) across all manage storage systems.

### Procedure

1. Click  in the title bar.
2. Select the type of object (**Storage Group, Initiator, Host, Virtual Machine and ESXi Server**).
3. Depending on the object you are looking for, type the following:
  - **Storage Group**—Type all or part of the storage group name.

- **Initiator**—Type all or part of the initiator name.
  - **Host**—Type all or part of the host name.
  - **Virtual Machine**—Type all or part of the virtual machine name.
  - **ESXi Server**—Type all or part of the ESXi Server 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.


## Modifying server logging levels

This procedure explains how to set the severity level of the alerts to log in the debug log. Once set, will only log events with the specified severity.


### Procedure

1. Select **SUPPORT**.
2. In the **Actions** panel, select **MODIFY SERVER LOGGING**.
3. Select a **Server Logging** level (WARN, INFO or DEBUG) and click **OK**.

## Exiting the console

To exit the console, click  in the title bar, select **Sign Out** and click **OK** to confirm.

## 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

• <a href="#">Managing settings</a> .....	48
• <a href="#">Setting preferences</a> .....	49
• <a href="#">Backing up the database server</a> .....	50
• <a href="#">Viewing database backups</a> .....	51
• <a href="#">Deleting database backups</a> .....	51
• <a href="#">Alert settings</a> .....	51
• <a href="#">Server alerts</a> .....	66
• <a href="#">Security</a> .....	67
• <a href="#">Viewing user sessions</a> .....	80
• <a href="#">Roles and associated permissions</a> .....	80
• <a href="#">Link and launch</a> .....	83
• <a href="#">Managing Database Storage Analyzer (DSA) environment preferences</a> .....	85
• <a href="#">Managing data protection preferences</a> .....	85
• <a href="#">Viewing authentication authority information</a> .....	86
• <a href="#">Local User and Authorization operations</a> .....	87
• <a href="#">Link and Launch operations</a> .....	87
• <a href="#">Entering PIN number</a> .....	87
• <a href="#">Report operations</a> .....	87


# Managing settings

## Before you begin

- To perform this operation, you must be a StorageAdmin or higher.

This procedure explains how to manage system settings.

## Procedure

1. Select  to open the **Settings** panel.

The following categories of settings are displayed (and the Preferences settings are displayed by default - see [Setting preferences](#) on page 49):

- Preferences
- System and Licences
- Users and Groups
- Symmetrix Access Control
- Management
- Data Protection
- Performance
- Unisphere Databases
- DSA Environment
- Alerts

2. Click on one of the following categories to view or modify its settings.

- **Preferences**—[Setting preferences](#) on page 49
- **System and Licences > License Usage**—[Viewing license usage](#) on page 928
- **System and Licences > Solutions Enabler**—[Viewing host-based licenses](#) on page 928
- **System and Licences > Symmetrix Entitlements**—[Viewing Symmetrix entitlements](#) on page 927
- **Users and Groups > Authentication**—[Viewing authentication authorities](#) on page 69
- **Users and Groups > Local Users**—[Viewing local users](#) on page 78
- **Users and Groups > User Sessions**—[Viewing user sessions](#) on page 80
- **Users and Groups > Authorized Users and Groups**—[Viewing the authorized users and groups list](#) on page 75
- **Symmetrix Access Control > Access Control Entries**—[Viewing access control entry details](#) on page 939
- **Symmetrix Access Control > Access Groups**—[Viewing access group details](#) on page 934
- **Symmetrix Access Control > Access Pools**—[Viewing access pools](#) on page 936



- **Management > Symmetrix Attributes**—[Setting system attributes](#) on page 874
- **Management > Link and Launch**—[Viewing link and launch client registrations](#) on page 84
- **Data Protection**—[Managing data protection preferences](#) on page 85
- **Performance > System Registrations**—[Viewing system registrations](#) on page 591
- **Performance > Dashboard Catalog**—[Managing dashboard catalog](#) on page 594
- **Performance > Real Time Traces**—[Viewing Real Time traces](#) on page 586
- **Performance > Metrics**—[Viewing and managing metrics](#) on page 718
- **Performance > Import Settings**—[Importing Performance settings](#) on page 604
- **Performance > Export Settings**—[Exporting Performance settings](#) on page 605
- **Performance > Export PV Settings**—[Exporting Performance Viewer settings](#) on page 605
- **Unisphere Databases > Performance Databases**—[Viewing Performance databases](#) on page 594
- **Unisphere Databases > System Database**—[Viewing database backups](#) on page 51
- **DSA Environment**—[Managing Database Storage Analyzer \(DSA\) environment preferences](#) on page 85
- **Alerts > Alert Policies**—[Configuring alert policies](#) on page 56
- **Alerts > Compliance Alert Policies**—[Viewing compliance alerts policies](#) on page 65
- **Alerts > Performance Thresholds and Alerts**—[Viewing Performance thresholds and alerts](#) on page 62
- **Alerts > Symmetrix Thresholds and Alerts**—[Viewing threshold alerts](#) on page 58
- **Alerts > Notifications**—[Configuring alert notifications](#) on page 55


## Setting preferences

### Before you begin

Only a user with Administrator permission can set preferences.

To set system preferences:

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Preferences** to open the **Preferences** page.
3. Modify any number of the following preferences:

**Unisphere 360 Support**—This setting enables or disables Unisphere 360 integration from the Unisphere side. Setting the checkbox to "disabled"

prevents Unisphere 360 from being able to enroll this Unisphere and disconnects any instance of Unisphere 360 that had previously enrolled it.

**Initial Setup User Warning**—This setting enables or disables the display of the warning when permissions are not configured during initial setup.

**Introduction to Health Score Card**—This setting enables or disables the display of health score guide in the **System Health** dashboard.

#### Custom Welcome Screen Message

Type a message to display to users during login. For example, you may want to notify logging in users about a software upgrade. Messages can be up to 240 characters.

#### Solutions Enabler Debug

Specify the 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.

---

#### Note

Changing the debug level from the default value of 0 might substantially increase the size of the log files and affect your system's performance.

---

4. Click **APPLY**.


## Backing up the database server

### Before you begin

To perform this operation, you must be an Administrator.

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 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.

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Unisphere Databases > System Database**
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 consists of a time stamp and your custom description.

5. Click **OK**.

## Viewing database backups

### Before you begin

To perform this operation, you must be a Monitor.


### Procedure

1. Select  to open the **Settings** panel.
2. Select **Unisphere Databases > System Database**

The following properties display:

- **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.
- **Description**—Message related to the backup.

The following controls are available:



- **Backup**—[Backing up the database server](#) on page 50
-  —[Deleting database backups](#) on page 51

## Deleting database backups

### Before you begin

To perform this operation, you must be an Administrator.

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Unisphere Databases > System Database**
3. Select one or more backups and click  .
4. Click **OK**.

## Alert settings

You can configure Unisphere to monitor storage systems for specific events or error conditions. When an event or error of interest occurs, Unisphere displays an alert and, if configured to do so, notifies you of the alert by way of email, SNMP, or Syslog.

In addition to alerting you of specific events or errors, Unisphere also generates a number of server alerts that also alert you. For more information, refer to [Server alerts](#) on page 66.

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

## Alerts

### Viewing alerts

#### Before you begin

- For alert (event) descriptions, refer to the [Events and Alerts Guide](#).
- In addition to alerting you of specific events or errors, Unisphere also generates a number of server alerts that also alert you. For more information, refer to [Server alerts](#) on page 66.
- The maximum number of alerts Unisphere displays is 10,000. Once this threshold is reached, Unisphere deletes the oldest alert for each subsequent alert it receives.

This procedure explains how to view alerts for a particular storage system or all the visible storage systems. This procedure also applies to storage container alerts which can be viewed by navigating to **STORAGE > VVol dashboard** and clicking on **STORAGE CONTAINER ALERTS** from within the **Actions** panel.

#### Procedure

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 **EVENTS > Alerts** to open the system's **Alerts** list view.

For all visible storage systems:

- a.  Select **Home**, and then select  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](#) on page 53.

In both cases, the following properties display:

- **State**—State of the alert. Possible values are New or Acknowledged.
- **Severity**—Severity of the alert. 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 storage 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 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:

-  —[Viewing alert details](#) on page 54.
- **Acknowledge**—[Acknowledging alerts](#) on page 53.
-  —[Deleting alerts](#) on page 54.


## Filtering alerts

### Procedure

1. Select **EVENTS** > **Alerts**, or select **Home** and then select  to open the Alerts list view.
2. Use the filter tool  to narrow the listed alerts to only those that meet the specified criteria:
  - **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.
  - **Symmetrix**— Filters the list based on the storage system identity.
  - **Object**—Filters the list for alerts for the specified object.
  - **Description**—Filters the list for alerts with the specified description.
  - **Created**— Filters the list based on when the alert was created.
  - **Acknowledged**— Filters the list for alerts that have been acknowledged.



## Acknowledging alerts

### Procedure

1. Select **EVENTS** > **Alerts**, or select **Home** and then select  to open the Alerts list view.
2. Select one or more alerts and click **Acknowledge**.

## Viewing alert details

### Procedure

1. Select **EVENTS** > **Alerts**, or select **Home** and then select  to open the Alerts list view.
2. Select an alert and click  to open the Alerts Details view.

The following properties display:

#### Alert ID

Unique number assigned by Unisphere.

#### 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.


#### Acknowledged

Shows the date on which the alert was acknowledged (if it has been).

## Deleting alerts

### Procedure

1. Select **EVENTS** > **Alerts**, or select **Home** and then select  to open the Alerts list view.

2. Select one or more alerts and click .

## Configuring alert notifications

### Before you begin

- To perform this operation, you must be an Administrator or StorageAdmin.
- Unisphere 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 system event exceeds a set threshold, the alert infrastructure ignores subsequent alerts from the source.

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

### Procedure

1. Do one of the following:
  - To enable alert notifications:
    - Select  to open the **Settings** panel.
    - Select **Alerts > Notifications** to open the **Notifications** page.
    - In the panel, click **Configure** for method you want to use to deliver the notifications (see [Configuring email notifications](#) on page 59 or [Configuring SNMP notifications](#) on page 602. (Not applicable for Syslog. For syslog, refer to *Setting up the event daemon for monitoring* in the *Solutions Enabler Installation and Configuration Guide* for instructions. )
    - In the panel, move the slider bar to the right to enable the configured 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.

- In the **Alerts** panel, do the following for each storage system from which you want to receive notifications:
  - Select the **System Level** and **Performance Level** severities in which you are interested.
  - To clear your selection, click a previously clicked item.

- Once satisfied, click **APPLY**.


## Alert policies

### Configuring alert policies

#### Before you begin

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

#### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Alert Policies**.
3. Select all or a specific storage system from the **Select Array** drop-down list.

The following properties display:

##### Name

Policy name. For alert (event) descriptions, refer to the *Solutions Enabler Installation Guide*.

##### Type

Type of alert policy. Possible values are:

- **Array** for array-based alerts.
- **SMAS** for application-based alerts.
- **File** for eNAS-based alerts.

##### Enabled

Whether the policy is Enabled or Disabled.

##### Notifications

Icon indicating the method to use when delivering the alert notification (e-mail, SNMP, or Sys Log). None indicates that Unisphere is not configured to deliver an alert notification for the corresponding policy.

4. To enable alert reporting for a particular event, [configure alert notifications](#), select the **Enabled** checkbox for that event and click **APPLY**.
5. To disable alert reporting for a particular event, clear the **Enabled** checkbox for that event and click **APPLY**.

## Threshold alerts

### Managing threshold alerts

#### Before you begin


- For alert (event) descriptions, refer to the [Events and Alerts Guide](#).
- Pool utilization thresholds are enabled by default on every storage 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.

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.

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Symmetrix Threshold and Alerts**.
3. Do the following, depending on whether you are creating, editing, or deleting thresholds:
  - Creating:
    - Click **Create**.
    - Select the storage system on which to create the threshold.
    - 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 5876.

#### DSE Spill Duration

Threshold event that reflects how long (in minutes) an SRDF spillover has been occurring. This category only applies to Enginuity 5876.

#### Snap Pool Utilization

Threshold event that reflects the allocated capacity (as percentage) within a snap pool. This category only applies to Enginuity 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 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.

#### Local Replication Utilization

Threshold event that indicates that the local replication resource usage has exceeded the threshold. This category only applies to storage systems running HYPERMAX OS 5977 Q1 2016 SR or higher.

#### System Meta Data Utilization

Threshold event that indicates that the system meta data Utilization usage has exceeded the threshold. This category only applies to storage systems running HYPERMAX OS 5977 Q1 2017 SR or higher.

### Storage Container Utilization





Threshold event that indicates that the storage container utilization usage has exceeded the threshold. This category only applies to storage systems running HYPERMAX OS 5977 or higher.

### Frontend Meta Data Usage

Threshold event that indicates that the front end meta data usage has exceeded the threshold. This category only applies to storage systems running PowerMaxOS 5978 or higher.

### Backend Meta Data Usage

Threshold event that indicates that the back end meta data usage has exceeded the threshold. This category only applies to storage systems running PowerMaxOS 5978 or higher.



- Select the pools (**Instances to enable**) on which to create the threshold.
- Enable (select) or disable (clear) alerts for the threshold.
- Specify a threshold value (percentage of utilization) for each severity level: **Warning**, **Critical**, and **Fatal**.
- Click **OK**.
- Editing:
  -  Hover over a threshold and click .
  - Select a threshold and specify a new threshold value (percentage of utilization) for any number of the severity levels: **Warning**, **Critical**, and **Fatal**.
  - Enable (select) or disable (clear) alerts for the threshold.
  - Click **OK**.
- Deleting:
  -  Hover over a threshold and click .

## Viewing threshold alerts

### Before you begin

- For alert (event) descriptions, refer to the [Events and Alerts Guide](#).

### Procedure


1.  Select  to open the **Settings** panel.
2. Select **Alerts > Symmetrix Threshold and Alerts**
3. Select All or a specific storage system.

The following properties display:

- **Name**—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.
- **Local Replication Utilization**—Threshold event that indicates that the local replication resource usage has exceeded the threshold. This category only applies to storage systems running HYPERMAX OS 5977 Q1 2016 SR or higher.
- **System Meta Data Utilization**—Threshold event that indicates that the system meta data Utilization usage has exceeded the threshold. This category only applies to storage systems running HYPERMAX OS 5977 Q1 2017 SR or higher.
- **Storage Container Utilization**—Threshold event that indicates that the storage container utilization usage has exceeded the threshold. This category only applies to storage systems running HYPERMAX OS 5977 or higher.
- **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.
- **Warning**—Percentage of utilization at which point a warning alert is issued.
- **Critical**—Percentage of utilization at which point a critical alert is issued.
- **Fatal**—Percentage of utilization at which point a fatal alert is issued.
- **Custom**—Whether the policy has been customized.
- **Enabled**—Whether the policy is Enabled or Disabled.
- **Notifications**—Whether the alert notification option is enabled (Email, SYSLOG, or SNMP) or disabled (NONE) for the alert.

The following controls are available:



- **Create**—[Managing threshold alerts](#) on page 56
-  —[Managing threshold alerts](#) on page 56

## Configuring email notifications

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 notifications:

### Procedure



1. To set up email notification:
  - a.  Select  to open the **Settings** panel.
  - b. Click **Alerts > Notifications**.

- c. In the **Email** section, click **Configure**.
- d. In the **Outgoing Mail Server (SMTP)** section specify the following details:
  - **IP Address/Host**
  - **Server Port**
- e. In the **User Information** section, specify the **Sender E-mail Address**.
- f. In the **Recipients** section, click **Create** and specify the address you want to add.
- g. Select one or more system or performance level indicators or reports to enable email notifications for the relevant level of system or performance notifications.
- h. Click **APPLY**.

## Editing subscriptions

To edit a subscription:

### Procedure

1. Select  to open the **Settings** panel.
2. Click **Alerts > Notifications**
3. Select a storage system and click  to open the **Edit Subscriptions** dialog.
4. Tick one or more of the checkboxes (**System Notifications**, **Performance Notifications**, and **Reports**) and click **OK**.


## Performance thresholds and alerts

### Creating a performance threshold 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 Diagnostic view.

To create performance threshold alerts:

### Procedure

1. Select  to open the **Settings** panel.
  2. Select **Alerts > Performance Thresholds and Alerts**.
  3. Select a storage system.
  4. Select the category for which you want to create a threshold or alert.
  5. Click **Create**.
- The **Create Threshold and Alert** wizard displays.
6. Select the **Array**, **Category**, and **Metrics**.
  - 7.

Select Instances from the **Available Instances** list and click



to move them to the **Instances to Enable** list.

8. Add a value for **Warning Threshold** or **Warning Threshold** and **Critical Threshold**.
9. Click **NEXT**.
10. To add an alert for each configured threshold, complete the following steps:
  - a. Select **Enable Alert**.
  - b. For each threshold you are configuring, specify values for the following fields:

#### Severity

The following values are available:

- **Information**
- **Warning**
- **Critical**

#### Occurrence

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.

#### Samples

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.

- c. (Optional) If required, 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.



11. Click **OK**.

## Editing a performance threshold 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 threshold alerts:

#### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Performance Thresholds and Alerts**.
3. Navigate to the threshold alert to be edited by selecting the appropriate storage system and category.
4. Hover over an item from the table and click  .
5. Edit the settings.
6. 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:


### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Performance Thresholds and Alerts**.
3. Navigate to the threshold or alert to be edited by selecting the appropriate category in the **Category** section.
4. Select one or more rows and click .
5. Click **OK**.

## Viewing Performance thresholds and alerts

You can configure a warning threshold and a critical threshold value for each metric.

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Performance Thresholds and Alerts**.
3. Select All or a storage system.
4. 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.

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

The following properties display:

- **Name**—The metric name.
- **Warning**—The warning threshold.
- **Alert**—Indicates if a warning alert has been generated. The icon displayed corresponds to the alert type.
- **Critical**—The critical threshold.
- **Alert**—Indicates if a critical alert has been generated.. The icon displayed corresponds to the alert type.
- **KPI**—Indicates if the metric is a KPI.
- **Custom**—Indicates if a custom threshold or alert has been generated..

5. Click **APPLY**.

The following controls are available:

- **Create**- [Creating a performance threshold alert](#) on page 60

-  - [Deleting performance thresholds and alerts](#) on page 62.

## Service level alert policies

### Creating service level compliance alerts policies

This procedure explains how to configure Unisphere to alert you when the performance of a storage group, relative to its service level target, changes. Once configured, Unisphere 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 service level. If both calculated values fall within (under) the service level 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** Service level 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 compliance alerts policies](#) on page 65.

**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 a service level other than optimized.
  - Contain volumes other than gatekeepers.
  - Be in a masking view.
  - Not have a policy currently associated with it.

To create service level compliance alert policies:


**Procedure**

1. Select  to open the **Settings** panel.
2. Select **Alerts > Compliance Alert Policies**
3. Click **Create**.
4. Select the storage system on which the storage groups are located.
5. Select one or more storage groups and click **>**.
6. (Optional) By default, service level compliance policies are configured to generate alerts for all service level 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 service level targets.
  -  Marginal—Storage group performing below service level target.
  -  Stable—Storage group performing within the service level target.
7. Click **OK**.

**Editing compliance alerts policies****Before you begin**

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

**Procedure**

1. Select  to open the **Settings** panel.
2. Select **Alerts > Compliance Alert Policies**
3. Select the policy, and then select (enable) or clear (disable) any of the compliance states.  
  
Unisphere generates alerts only for enabled compliance states.
4. Click **APPLY**.





## Deleting compliance alerts policies

### Before you begin

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


### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Compliance Alert Policies**
3. Select one or more policies and click .
4. Click **OK**.




## Viewing compliance alerts policies

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


### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Compliance Alert Policies**
3. Select All or a specific storage system.

The following properties display:

- **Name** — Policy name..
- **Compliance State** — Enabled compliance states:
  -  Critical — Storage group performing well below service level targets
  -  Marginal — Storage group performing below service level target.
  -  Stable — Storage group performing within the service level target.
- **Notifications** — Method to use when delivering the alert notification (e-mail, SNMP, or Sys Log). None indicates that Unisphere is not configured to deliver an alert notification for the corresponding policy. To enable alert reporting for a particular event, see [Configuring compliance alert notifications](#) on page 66.

The following controls are available:

- **Create**—[Creating service level compliance alerts policies](#) on page 63
-  —[Deleting compliance alerts policies](#) on page 65


## Configuring compliance alert notifications

### 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 55 .

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

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Compliance Alert Policies**
3. Select one or more policies and click **Notify**.
4. Select (enable) the method you want to use to deliver the notifications.  
Possible methods are:
  - **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.

---

#### Note

The storage system must already be configured to deliver alerts in the desired method, as described in [Configuring alert notifications](#) on page 55.

---

5. Click **APPLY**.

## Server alerts

Unisphere generates server alerts under the conditions listed in the table below.

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 also apply to discover operations. That is, performing a discover operation will not force the delivery of these alerts.

---

#### Note

Runtime alerts are not storage system-specific. They can be deleted as long as the user has admin or storage admin rights on at least one storage system. A user with a monitor role is not allowed to delete the server alerts.

---

Server alert	Number of volumes	Threshold	Alert Details
Total memory on the Unisphere server	0 - 64,000	12 GB	System memory <# GB> is below the minimum requirement of <# GB>
	64,000 - 128,000	16 GB	
	128,000 - 256,000	20 GB	

Server alert	Number of volumes	Threshold	Alert Details
Free disk space on the Unisphere installed directory	0 - 64,000	100 GB	Free disk space <# GB> is below the minimum requirement of <# GB>
	64,000 - 128,000	140 GB	
	128,000 - 256,000	180 GB	
Number of managed storage systems	Threshold is 20.		Number of managed arrays <#> is over the maximum supported number of #
Number of managed volumes	256,000		Number of managed volumes <#> is over the maximum supported number of <#>. 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

### Authentication

#### Login authentication

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

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

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

#### Logging in

The **Login** dialog box contains the following elements:

**Username**—user name (refer to [Login authentication](#) on page 67).

**Password**—password.

This dialog box may also include a login message. The login message feature enables Administrators and StorageAdmins to display a message to users during login. For example, an administrator may want to notify users about a software upgrade.

**Login**—Opens the console.

## Configuring authentication authorities


### Before you begin

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

This procedure explains how to configure Unisphere to authenticate users.

To configure authentication:



### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Authentication**.
3. Select the **Authentication Authority** to use during login. Possible values are:
  - **Local Directory** — You can disable this if enabled and enable this if disabled. When enabled, users can log in as a user from the CST local directory.
  - **LDAP-SSL** — You can disable this if enabled and enable this if disabled. When enabled, users can log in as a user from the configured LDAP directory.
  - **Windows OS/AD** — You can disable this if enabled and enable this if disabled. 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.
4. If you select the Windows OS/AD authority and click **Modify**, as an option you can specify to limit authentication to members of a specific Windows OS/AD group. To do this, select the **Limit authentication to members of a specific Windows OS/AD group(s)** checkbox option, and type the **Group Name(s)**, separated by commas.
5. Click **Next**.
6. If you are configuring LDAP-SSL, click **Enable** and 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 **Choose File**, locate the certificate, and click **Open**. To view the contents of the certificate, click **VIEW CERTIFICATE**. To clear the file selection, click **CLEAR**.
- c. Optional: To limit authentication to only members of specific LDAP groups, click **Limit Authentication to members of LDAP group(s)**, select the option, and then type the **Group Name(s)**, separated by commas.
- d. Click **Next**.
7. Click **OK**.

## Viewing authentication authorities

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Settings > Users and Groups > Authentication**  
Use the **Authentication** page to view and manage authentication settings.  
The following properties display:  
**Authentication**  
The following authentication types are displayed:
  - **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.
 The following controls are available:
  -  —Hover over an authentication type and click to view the authentication authority information (see [Viewing authentication authority information](#) on page 86).

- **Enable**—This control changes the status of **Local Directory** or **Windows OS/AD** from disabled to enabled. This control also changes the status of **LDAP-SSL** (see [Configuring authentication authorities](#) on page 68 ).
- **Disable**—This control changes the status of **Local Directory**, **Windows OS/AD**, or **LDAP-SSL** from enabled to disabled.
- **Modify**—[Configuring authentication authorities](#) on page 68

## 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 80.

- **None**—Provides no permissions.
- **Monitor**—Performs read-only (passive) operations on a storage system excluding the ability to read the audit log or Access Control definitions.
- **StorageAdmin**—Performs all management (active or control) operations on a 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, (including reading the audit log, symacl list and symauth) in addition to all Monitor operations. This is the minimum role required to view the Symmetrix audit log.
- **DSA Admin**—Collects and analyzes database activity with Database Storage Analyzer.  
A user cannot change their own role so as to remove Administrator or SecurityAdmin privileges from themselves.
- **Local Replication**—Performs local replication operations (SnapVX or legacy Snapshot, Clone, BCV). To create Secure SnapVX snapshots a user needs to have

Storage Admin rights at the array level. This role also automatically includes Monitor rights.

- **Remote Replication**—Performs remote replication (SRDF) operations involving devices and pairs. Users can create, operate upon or delete SRDF device pairs but can't create, modify or delete SRDF groups. This role also automatically includes Monitor rights.
- **Device Management**—Grants user rights to perform control and configuration operations on devices. Storage Admin rights are required to create, expand or delete devices. This role also automatically includes Monitor rights.

In addition to these user roles, Unisphere includes an administrative role, the Initial Setup User. This user, defined during installation, is a temporary role that provides administrator-like permissions for the purpose of adding local users and roles to Unisphere. For more information, see [Operating as the initial setup user](#) on page 32.

## 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 (for example, `D:sales\putman`) is found, that is used; if a partial match (for example, `D:*putman`) is found, that is used; if an unqualified match (for example, `putman`) is found, that is used; otherwise, the user is assigned a role of None.

## 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:

#### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Authorized Users and Groups**
3. Click **Create**.
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	Unisphere server hostname
Windows OS	
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 **DSA Fast Hinting** role. This role is only enabled for DSA administrators.
  9. If you choose a **Local Replication**, **Remote Replication** or **Device Management** role, click **Select Storage Group(s)** and in the edit dialog that opens choose between:
    - a. **Wildcard**—A wildcard syntax used with the storage group component name to allow a single rule to apply to multiple storage groups.  
A simple wildcard syntax can be used with the component name to allow a single rule to apply to multiple SGs as follows:  
 abc - Exactly these characters  
 ? Any 1 character  
 \* Any zero or more characters  
 + Zero or more additional occurrences of the previous match  
 [a-z0-9] Any of these characters  
 [!a-z] Anything but one of these characters  
 All SG name comparisons are case-insensitive. The following examples show how they will be interpreted:

This pattern	Matches these Storage Groups	Does not match these Storage Groups
tg_*	tg_DB_SG1 or tg_newSG or TG_sg_db	tgNewSG
prod_sg?	prod_sg1 or prod_sga por Prod_sg2	prod_sg12 or prod_sgab
prod_sg[0-9]+	prod_sg1 or prod_sg12	prod_sga or prod_sgab

The only allowed characters are: a-zA-Z0-9\_- along with the above \*+?[!] wildcard characters.

The only roles that can be assigned against storage groups are: Local Replication, Remote Replication and Device Management.

Storage groups do not have to exist at the time that a matching Role Based Authentication Controls (RBAC) rule for them is defined.

These storage groups-level RBAC rules are only applicable to parent and stand-alone SGs and not child SGs. Child SGs are protected by the RBAC rules, if any, on their parent SG.

---

**Note**

Unisphere for PowerMax does not support RBAC Device Group management.

---

**b. Storage Group**

c. Once your input or selection is complete, click **Save**.

10. Click **OK**.


## Editing 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 storage systems.

To modify authorization rules:

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Authorized Users and Groups**
3. Select a storage system ID from the drop-down list.
4. Select a rule and click **Modify**.
5. On the **Roles** tab, add or remove from any of the available objects, being sure to not exceed the four roles/object limit.
6. 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 on all authorized storage systems.

---



**Note**

To remove an authorization rule on a single object, select **Modify**.

---


To remove authorization rules:

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Authorized Users and Groups**
3. Select a storage system ID from the drop-down list.
4. Select a rule and click .
5. Click **OK**.

## Viewing authorization rules

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Authorized Users & Groups**.

### After you finish


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.
  - **Unsupported**—Not supported.
- **Authentication Domain**—Domain name. Possible values are based on the authentication authority:

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

The following controls are available:


- **Create**—[Adding authorization rules](#) on page 72
- **Remove**  —[Removing authorization rules](#) on page 74
- **Modify**—[Editing authorization rules](#) on page 74


## Viewing the authorized users and groups list

To view local user details, refer to [Viewing local users details](#) on page 78.


To view the authorized users and groups list:

### Procedure

1. Select  to open the **Settings** panel.

2. Select **Users and Groups > Authorized Users & Groups**.
3. Select your required storage system ID from the drop-down list.
4. To see more information on a user, select the user and on the right-hand side of the row, click the  icon.



The following controls are available:

- **Create**- [Adding authorization rules](#) on page 72
- **Modify**- [Editing authorization rules](#) on page 74
-  - [Removing authorization rules](#) on page 74

## Viewing the authorized users and groups details

To view the authorized users and groups list:

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Authorized Users & Groups**.
3. Select your required storage system ID from the drop-down list.
4. To see more information on a user, select the user and on the right-hand side of the row, click the  icon.
5. View the following information in the information dialog: name, authority, domain, storage system identity, roles, and component name.

## View Certificate dialog box

Use this dialog box to view contents of an SSL certificate.

## Local Users


### Creating local users

#### Before you begin

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

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

#### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Local Users**
3. Click **Create** to open the **Create Local User** 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. There are no restrictions on special characters when using passwords. However, these characters should not be used when creating user names: \ : ,
7. Select the storage system and click the **Roles** tab and select one or more roles - up to four can be selected.
8. For **Database Storage Analyzer** select **None**, **Read Only**, or **Admin**.  
 By default, **Database Storage Analyzer** permissions are set to **Read Only**.
9. Click the **DSA Fast Hinting** role.  
 This role is only enabled for DSA administrators. It allows a user to create and modify FAST hints.
10. Click **OK**.

## Editing local users



### Before you begin

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

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

To create local users:

### Procedure

1.  Select  to open the **Settings** panel.
2. Select **Users and Groups > Local Users**
3. Select a user and click **Modify**.
4. Optional: Type a new **Description**.
5. On the **Roles** tab, add or remove from any of the available objects, being sure to not exceed the four roles/object limit.
6. Click **OK**.




## Deleting local users

### Before you begin

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

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

### Procedure

1.  Select  to open the **Settings** panel.
2. Select **Users and Groups > Local Users**
3. Select a user and click  .

4. Click **OK**.

## Changing 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 storage system.

This procedure explains how to change a local user passwords.

To change local directory user passwords:

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Local Users**
3. Select a user and click **Change Password**.
4. Type the user's **Old Password**.
5. Type a **New Password** and **Confirm Password**.
6. 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 storage system.

To view users with a local Unisphere account:



### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Local Users**
3. 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:

-  —[Viewing local users details](#) on page 78.
- **Create**—[Creating local users](#) on page 76.
- **Modify**—[Editing local users](#) on page 77.
- **Change Password**—[Changing local user passwords](#) on page 78.
-  —[Deleting local users](#) on page 77.

## Viewing local users details

This procedure explains how to view the details of a local user.

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Local Users**.
3. Select a user, hover over the row and click  to see the details view.



The following properties display:

- **Name**—User or group name.
- **Authority**
- **Domain**
- **Symmetrix ID**
- **Roles**
- **Component Name**

## Viewing authorization rules

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


### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Authorized Users & Groups**.
3. Select the user and click  to open the user's details view.
4. 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	Unisphere server hostname
Windows OS	
Windows AD	Unisphere server domain

Authority	Domain name
LDAP-SSL	LDAP server domain
Virtualization domain	Virtualization domain
Any authority	Any


The following controls are available:

- **Create**—[Adding authorization rules](#) on page 72
- **Modify**—[Editing authorization rules](#) on page 74
- **Delete**  —[Removing authorization rules](#) on page 74

## Viewing user sessions

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

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > User Sessions**

The following properties display:

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

## Roles and associated permissions

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

### Note

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

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

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



- Local Replication
- Remote Replication
- Device Management

**Table 2** User roles and associated permissions

Permissions	AD	SA	MO	SecA	AUD	None	PM	DSA
Create/delete user accounts	Yes	No	No	Yes	No	No	No	No
Reset user password	Yes	No	No	Yes	No	No	No	No
Create roles	Yes	Yes	No	Yes (self-excluded)	No	No	No	No
Change own password	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Manage storage systems	Yes	Yes	No	No	No	No	No	No
Discover storage systems	Yes	No	No	Yes	No	No	No	No
Add/show license keys	Yes	Yes	No	No	No	No	No	No
Set alerts and Optimizer monitoring options	Yes	Yes	No	No	No	No	No	No
Release storage system locks	Yes	Yes	No	No	No	No	No	No
Set Access Controls	Yes	Yes	No	No	No	No	No	No
Set replication and reservation preferences	Yes	Yes	No	No	No	No	No	No
View the storage system audit log	Yes	Yes	No	Yes	Yes	No	No	No
Access performance data	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Start data traces	Yes	Yes	Yes	Yes	Yes	No	Yes	No
Set performance thresholds/alerts	Yes	Yes	No	No	No	No	Yes	No
Create and manage	Yes	Yes	Yes	Yes	Yes	No	Yes	No

**Table 2** User roles and associated permissions (continued)

Permissions	AD	SA	MO	SecA	AUD	None	PM	DSA
performance dashboards								
Collect and analyze database activity with DSA	No	No	No	No	No	No	No	Yes

**Table 3** Permissions for Local Replication, Remote Replication and Device Management roles

Permissions	Local Replication	Remote Replication	Device Management
Create/delete user accounts	No	No	No
Reset user password	No	No	No
Create roles	No	No	No
Change own password	Yes	Yes	Yes
Manage storage systems	No	No	No
Discover storage systems	No	No	No
Add/show license keys	No	No	No
Set alerts and Optimizer monitoring options	No	No	No
Release storage system locks	No	No	No
Set Access Controls	No	No	No
Set replication and reservation preferences	No	No	No
View the storage system audit log	No	No	No
Access performance data	Yes	Yes	Yes
Start data traces	Yes	Yes	Yes
Set performance thresholds/alerts	No	No	No
Create and manage performance dashboards	Yes	Yes	Yes
Collect and analyze database activity with	No	No	No

**Table 3** Permissions for Local Replication, Remote Replication and Device Management roles (continued)

Permissions			
	Local Replication	Remote Replication	Device Management
Database Storage Analyzer			
Perform control, configuration and expand operations on devices	No	No	Yes
Create or delete devices	No	No	No
Perform local replication operations (SnapVX, legacy Snapshot, Clone, BCV)	Yes	No	No
Create Secure SnapVX snapshots	Yes	No	No
Create, operate upon or delete SRDF device pairs	No	Yes	No
Create, modify or delete SRDF groups	No	No	No

## Link and launch

### Creating link-and-launch client registrations


#### Before you begin

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

Link-and-launch is not supported with X.509 certificate-based user authentication.

This procedure explains how to register other applications with the SMAS server. Once registered, users of the registered applications can launch Unisphere without logging in.

#### Procedure

1. Select  to open the **Settings** panel.
2. Select **Management > Link and Launch**
3. Click **Create**.
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

### Before you begin

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

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

### Procedure



1. Select  to open the **Settings** panel.
2. Select **Management > Link and Launch**
3. Select a registration, and click **Edit**.
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

### Before you begin


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

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Management > Link and Launch**
3. Select a registration, and click .
4. Click **OK**.

## Viewing link and launch client registrations

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Management > Link and Launch**


### After you finish

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

The following property displays:

- **Client ID**—Unique client ID.

The following controls are available:

- **Create**—[Creating link-and-launch client registrations](#) on page 83
- **Edit**—[Editing link-and-launch client registrations](#) on page 84
-  —[Deleting link-and-launch client registrations](#) on page 84



## Managing Database Storage Analyzer (DSA) environment preferences

### Before you begin

Only a user with Administrator permission can specify DSA environment preferences.

To specify DSA environment preferences:

### Procedure

1.  Select  to open the **Settings** panel.
2. Select **DSA Environment** to open the **DSA Environments** page.
3. Select an environment from the Environments drop-down list.
4. Modify any number of the following:
  - **Data retention for 5 min data**—Number of days (between 15 and 30) to retain 5 minute data.
  - **Data retention for hourly data**—Number of months (between 12 and 36) to retain hourly data.
  - **Data retention for daily data**—Number of months (between 12 and 36) to retain daily data.
  - **First threshold for DB read response time (ms)**—First threshold for DB read response time.
  - **Second threshold for DB read response time (ms)**—Second threshold for DB read response time.
5. Click **APPLY**.



## Managing data protection preferences

### Before you begin

Only a user with Administrator permission can specify data protection preferences.

To specify data protection preferences.:



### Procedure

1.  Select  to open the **Settings** panel.
2. Select **Data Protection** to open the **Data Protection** page.
3. Modify any number of the following:
  - **Clone Copy Mode**—Select the default behavior for creating clone sessions. Possible values are:
    - **No Copy No Diff**—Creates 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 are 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 are copied).
  - **VSE No Diff**—Creates a VP Snap Copy session.
  - **Clone Target**—Select the default target volume.
  - **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.
4. Click **APPLY**.

## Viewing authentication authority information

### Procedure

1. Select  to open the **Settings** panel.
  2. Select **Settings > Users and Groups > Authentication**
  3. Hover over an authentication type and click .
- The authentication authority information is displayed. For LDAP-SSL, the following is displayed when LDAP is enabled.)
- **Port**—Port number of the LDAP service. Typically, this value is 389 for LDAP and 636 for LDAPS.
  - **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.
- Status**—Status of authentication (enabled or disabled).
- Limited Authentication Group(s)**—Limited Authentication Group(s) names.

## Local User and Authorization operations

- Modify Local User (see [Editing local users](#) on page 77).
- Viewing local users details (see [Viewing local users details](#) on page 78).
- Modify Authorization rules (see [Removing authorization rules](#) on page 74).
- Viewing the authorized users and groups details (see [Viewing the authorized users and groups details](#) on page 76).


## Link and Launch operations

- Creating link-and-launch client registrations (see [Creating link-and-launch client registrations](#) on page 83).
- Editing link-and-launch client registrations (see [Editing link-and-launch client registrations](#) on page 84).

## Entering PIN number

To enter the PIN number :

### Procedure

1. Select  to open the **Settings** panel.
2. Select one of the following: .
  - **Symmetrix Access Control > Access Control Entries**
  - **Symmetrix Access Control > Access Groups**
  - **Symmetrix Access Control > Access Pools**

A warning is displayed if you have read- only access.
3. Click **Enter PIN**.
4. Enter the PIN number.
5. Click **OK**.

## Report operations

The following report operations are available:

- Create Compliance reports (see [Creating Compliance Reports](#) on page 157).
- Create performance reports (see [Creating performance reports](#) on page 579).

- Modify performance reports (see [Modifying performance reports](#) on page 583).
- Copy performance reports (see [Copying performance reports](#) on page 580).



# CHAPTER 4

## Storage Management

• Understanding Storage Management.....	90
• Tag and Untag operations.....	91
• Viewing Storage Group Demand Reports.....	91
• Viewing Service Level Demand Reports.....	92
• Viewing CKD volumes.....	92
• Viewing CKD volumes in CU image.....	93
• Viewing Storage Group Compliance view.....	94
• Dialog displayed when there is less than one week's data collected.....	96
• Setting volume emulation.....	96
• FAST association operations.....	97
• Removing DATA volumes.....	97
• Mapping volume operations.....	97
• Rename operations.....	98
• Provisioning storage .....	98
• Creating storage groups.....	112
• Understanding FAST .....	151
• Managing volumes.....	177
• Viewing disk groups.....	230
• Creating DATA volumes.....	237
• Creating thin pools.....	240
• Creating thin volumes.....	256
• Understanding Virtual LUN Migration.....	258
• Understanding Federated Tiered Storage .....	263
• Understanding storage templates.....	267
• Understanding FAST.X.....	272
• Viewing reservations.....	277
• Managing vVol.....	278
• Understanding compression.....	287

# Understanding Storage Management

Storage Management covers the following areas:

- **Storage Groups** - Management of Storage groups. Storage groups are a collection of devices stored on the array that are used by an application, a server, or a collection of servers. Storage groups are used to present storage to hosts in masking/mapping, Virtual LUN Technology, FAST, and various base operations.
- **Service Levels**- Management of service levels. A service level is the response time target for a storage group. The Service Level sets the storage array with the desired response time target for a storage group. It automatically monitors and adapts to the workload in order to maintain the response time target. The Service Level includes an optional workload type so it can be fine tuned to meet performance levels..
- **Templates** - Management of templates. Using the configuration and performance characteristics of an existing storage group as a starting point, you can create templates that will pre-populate fields in the provisioning wizard and create a more realistic performance reservation in your future provisioning requests.
- **Storage Resource Pools** - Management of Storage Resource Pools. Fully Automated Storage Tiering (FAST) provides automated management of storage array disk resources to achieve expected service levels. FAST automatically configures disk groups to form a Storage Resource Pool (SRP) by creating thin pools according to each individual disk technology, capacity and RAID type.
- **Volumes** - Management of volumes. A storage volume is an identifiable unit of data storage. Storage groups are sets of volumes.
- **External Storage** - Management of external storage. FAST.X attaches external storage to storage systems directs workload movement to these external arrays while having access to the array features such as local replication, remote replication, storage tiering, data management, and data migration. In addition, it simplifies multi-vendor or Dell EMC storage array management.
- **Vvol** - Management of Vvols. VMware VVols allow data replication, snapshots, encryption and so on to be controlled at the VMDK level instead of the LUN level, where these data services are performed on a per VM (application level) basis from the storage array.
- **FAST Policies** - Management of 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.
- **Tiers** - Management of storage tiers. FAST automatically moves active data to high-performance storage tiers and inactive data to low-cost, high-capacity storage tiers.
- **Thin Pools** - Management of Thin pools. Storage systems are pre-configured at the factory with virtually provisioned devices. Thin Provisioning helps reduce cost, improve capacity utilization, and simplify storage management. Thin Provisioning presents a large amount of capacity to a host and then consumes space only as needed from a shared pool. Thin Provisioning ensures that thin pools can expand in small increments while protecting performance, as well as non-disruptive shrinking of thin pools to help reuse space and improve capacity utilization.

- **Disk Groups** - Management of disk groups. A disk group is a collection of physical drives within the storage array that share the same performance characteristics.
- **Vlun Migration** - Management of VLUN migration. Virtual LUN Migration (VLUN Migration) enables transparent, nondisruptive data mobility for both disk group provisioned and virtually provisioned storage system volumes between storage tiers and between RAID protection schemes. Virtual LUN can be used to populate newly added drives or move volumes between high performance and high capacity drives, thereby delivering tiered storage capabilities within a single storage system. Migrations are performed while providing constant data availability and protection.

## Tag and Untag operations

The following tag and untag operations are available:

- **Storage Group level - RecoverPoint tag and untag** (see [Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472).
- **Volume level - RecoverPoint tag and untag** (see [Tagging and untagging volumes for RecoverPoint \(volume level\)](#) on page 472).
- **Data Protection> Open Replicator> RecoverPoint volumes untag** (see [Untagging RecoverPoint tagged volumes](#) on page 473).

## Viewing Storage Group Demand Reports

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

Before you begin:

This feature requires HYPERMAX OS 5977 or higher.

To view storage group demand reports:

### Procedure

1. Select the storage system.
2. Select **CAPACITY** to open the **CAPACITY** dashboard.
3. Select a SRP instance from the drop down menu and in the **Actions** panel, click **STORAGE GROUP DEMAND**.

Some or all of the following properties display:

- **Storage Group**—Name of the storage group.
- **Subscription (GB)**—Amount of SRP capacity to which the storage group subscribed.
- **Allocated (GB)**—The amount of allocated pool capacity (in GB).
- **Allocated (%)**—The percentage of allocated pool capacity.
- **Used (GB)**—The amount of allocated pool capacity (in GB).
- **Snapshot Allocated (GB)**—The amount allocated to snapshots.
- **Compression Ratio**—The compression ratio.
- **SNAP Used (GB)**—The amount used by snapshots.
- **Snapshot Compression Ratio**—The snapshot compression ratio.

- **Emulation**—Emulation type. This displays only if the storage system is capable of containing CKD devices.

The following control is available:

- **Export Report**—Exports the report to a PDF file.

## Viewing Service Level Demand Reports

This procedure explains how to view demand that each service level is placing on the SRP.

Before you begin:

This feature requires HYPERMAX OS 5977 or higher.

To view service level demand reports:

### Procedure

1. Select the storage system.
2. Select **CAPACITY** to open the **CAPACITY** dashboard.
3. Select a SRP instance from the drop down menu and in the **Actions** panel, click **SERVICE LEVEL DEMAND**.

### Results

Some or all of the following properties display:

#### Service Level

Name of the service level.

#### Allocated (GB)

Total space that the service level has allocated on the SRP in GB.

#### Allocated (%)

Percentage of space that the service level has allocated on the SRP.

#### Subscription (GB)

Total space that the service level has subscribed on the SRP in GB.

#### Subscription (%)

Percentage of space that the service level has subscribed on the SRP.

## Viewing CKD volumes

See below for procedure to view CKD volumes from the **Hosts > Mainframe** dashboard. To see the CKD volumes in a CU image, see [Viewing CKD volumes in CU image](#) on page 93.

### Procedure



1. Select the storage system.
2. Select **HOSTS > Mainframe** and click on **CKD Volumes** in the **Summary** panel.

The **CKD Volumes** list view is displayed. Use the this list view to view and manage the volumes.

The following properties display: however, not all properties may be available for every volume type:

- **Name**—Assigned volume name.
- **Type**—Type of volume.
- **Allocated %**—% of the volume that is allocated.
- **Capacity (GB)**—Volume capacity in Gigabytes.
- **Status**—Volume status.
- **Emulation**—Emulation type for the volume.
- **Host Paths**—Number of masking records for the volume.
- **Reserved**—Indicates whether the volume is reserved.
- **Split**—The name of the associated split.
- **CU Image** —The number of the associated CU image.
- **Base Address**—Base Address.


The following controls are available, however, not all controls may be available for every volume type:

-  —[Viewing CKD volume details](#) on page 204
- **Create**—[Creating volumes](#) on page 178
- **Expand**—[Expanding existing volumes](#) on page 191
-  —[Deleting volumes](#) on page 188
- **Create SG**—[Creating storage groups](#) on page 112
- **Set Volumes > Emulation**—[Setting volume emulation](#) on page 96
- **Set Volumes > Attribute**—[Setting volume attributes](#) on page 195
- **Set Volumes > Identifier**—[Setting volume identifiers](#) on page 196
- **Set Volumes > Status**—[Setting volume status](#) on page 194
- **Set Volumes > Replication QoS**—[QoS for replication](#) on page 197
- **Set Volumes > Set SRDF GCM**—[Setting the SRDF GCM flag](#) on page 434
- **Set Volumes > Reset SRDF/Metro Identity**—[Resetting original device identity](#) on page 432
- **Allocate/Free/Reclaim > Start** —[Managing thin pool allocations](#) on page 244
- **Allocate/Free/Reclaim > Stop**— [Managing thin pool allocations](#) on page 244
- **Configuration > Change Volume Configuration**—[Changing volume configuration](#) on page 190
- **Configuration > Duplicate Volume**—[Duplicating volumes](#) on page 188
- **Configuration > z/OS Map**—[z/OS map from the volume list view](#) on page 333
- **Configuration > z/OS Unmap**—[z/OS unmap from the volume list view](#) on page 334

## Viewing CKD volumes in CU image

Viewing CKD volumes in CU image

**Procedure**

1. Select the storage system.
2. Select **Hosts > CU Images**
3. Select the CU image and click .
4. In the details panel, click on the number in the **Number of Volumes** field to open the **CKD Volumes** list view..
5. Use the **CKD Volumes** list view to display and manage CKD volumes in a CU image.

**Results**

**Name** — Symmetrix volume name.

**Type** — Volume configuration.

**Status** — Volume status.

**Capacity (GB)** — Volume capacity in GBs.

**Emulation** — Emulation type.

**UCB Address** — Unit control block (address used by z/OS to access this volume).

**Volser** — Volume serial number (disk label (VOL1) used when the volume was initialized).

The following controls are available:



— [Viewing CU image details](#) on page 329

**z/OS Map**—[z/OS map from the volume list view](#) on page 333

**z/OS Unmap**—[z/OS unmap from the volume list view](#) on page 334

## Viewing Storage Group Compliance view

**Before you begin**

The user requires a minimum of Monitor permissions to perform this task.


**Definitions:**

- **Workload Skew** - Skew is represented by capacity and load pairs. There are two sources of skew for a storage group. One is using device stats. The other is using SG\_PER\_POOL chunks. There is an algorithm in WLP to merge these two lists to give us a usable skew profile. A skew profile is only useful if you have multiple chunks. If an SG has a single device, there is not enough data to calculate skew, the corresponding storage group per pool metrics can be used. Similarly, if an array has only one pool, the device stats are more meaningful for skew.
- **Workload Mixture** - The mixture is the distribution of various I/O types as percentages of the total IOPS. These are useful for determining, for example, whether a workload is heavy read or heavy write, whether I/Os are mostly random or mostly sequential.

To view the Storage Group (SG) Compliance view:

**Procedure**

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Groups View**.

3. Select a storage group and click  to view its details.
4. Select **VIEW ALL DETAILS**.
5. Select the **Compliance** tab.

Charts are displayed for the following:

- Response Time chart - this chart displays wait time weighted response time and (if applicable) the target service level response time band. The following section explains the data in the chart.
  - Actual: running I/O to Storage Group - Wait time weighted response time is calculated in buckets and displayed. If a bucket has no data, 0 is displayed.
  - Actual: no I/O to Storage Group - 0s are displayed.
  - Planned: SLO Response Time Max and SLO Response Time Min are displayed as a data band across the timeline. This is labeled "Planned". If the service level is Optimized, no plan is displayed, because there is no Response Time band for Optimized.
  - Excluded Data: If a recurring exclusion has been set via the Exclusion Windows dialog, the windows are represented by vertical gray plot bands.
  - Last Processed: A 2px dotted plot line marks the most recent SPA HOURLY timestamp processed by SPA for a given metric. It is not represented in the legend, but if you hover, you can see the timestamp associated. In normal successful/processing, this acts as a "Where am I" indicator. If WLP stops processing for some reason, it's a subtle debugging helper.
- IOPS chart - This chart toggles between IO/sec and MB/sec, displaying IO rate weighted metric values, "planned" values, and (if set) Host IO Limits. The following section explains the data in the chart.
  - Actual: running I/O to Storage Group - IO Rate weighted total IOPS (or total MBPS) are calculated in buckets and displayed. If a bucket has no data, 0 is displayed.
  - Actual: no I/O to Storage Group - 0s are displayed.
  - Planned: Host I/O Limits for Standalone SG - Host IO Limit is displayed as a static value across the timeline. Host IO Limit is only shown on the chart it impacts. For example, if MBPS host IO limit is set, and the user has IOPS selected, they won't see anything unless they toggle to MBPS.
  - Planned: Host I/O Limits for Child SG, no limit for the parent SG - Host IO Limit is displayed as a static value across the timeline. Host IO Limit is only shown on the chart it impacts. For example, if MBPS host IO limit is set, and the user has IOPS selected, they won't see anything unless they toggle to MBPS.
  - Planned: No Host I/O Limits for Child SG and parent SG - If a cascaded SG has a host IO limit set at the parent, but no direct limit of its own, the host IO limit of any given child would be the parent limit minus whatever the siblings are using.
  - Planned: Host I/O Limits for Child SG and parent SG - If a cascaded SG has a host IO limit set at the parent, and a direct limit of its own, the host IO limit of any given child would be the more limiting of the parent limit minus whatever the siblings are using, or the child SGs own limit.
  - Excluded Data: If a recurring exclusion has been set via the Exclusion Windows dialog, the windows are represented by vertical gray plot bands.

- Last Processed: A 2px dotted plot line marks the most recent SPA HOURLY timestamp processed by SPA for a given metric. It is not represented in the legend, but if you hover, you can see the timestamp associated. In normal successful/processing, this acts as a "Where am I" indicator. If WLP stops processing for some reason, it's a subtle debugging helper.
- Workload Skew chart - This chart compares actual workload skew - represented by cumulative capacity and load percentages (ordered by access density) - to planned skew. If there is no IO data, Actual is displayed as 50% skew - a straight line from (0,0) to (100,100). If there is one Device in SG AND Only One Thin Pool, then the merged device and sg per pool skew profile doesn't give us enough data points. Actual is displayed as 50% skew - a straight line from (0,0) to (100,100). If IO is running to the SG, the skew is a logarithmic curve (or stepped line graph in some cases).
- I/O Mixture chart - This chart compares actual workload mixture to planned workload mixture. The inner pie represents the actual IO distribution. The outer donut represents the planned mixture. If there is no I/O to the storage group, the mixture distribution will be equal percentages for each IO type (20% read hit, 20% sequential write, etc.) and the tooltip will show the corresponding IO sizes as 0kB.

Select the **Show Plan** slider to turn on or turn off the display of the plan. The plan is reference point used for comparison, and is a two week expiring performance reservation for subsequent provisioning suitability calculations.

The following controls are available:

- **Exclude Data** - [Managing Data Exclusion Windows](#) on page 158
- **Save As a Template** - [Creating storage templates](#) on page 267
- **Reset Workload Plan** - [Resetting Workload Plan](#) on page 177
- **Set Host I/O Limits** - [Setting host I/O limits](#) on page 132

## Dialog displayed when there is less than one week's data collected

This dialog is displayed when at least one week of data has not been collected for the selected storage group. It is recommended that you wait until you have at least one week of data. Alternatively, if you wish to proceed, click the **Autofill the template workload with averages from the stats collected so far** checkbox. Click **OK**.

## Setting volume emulation

### Before you begin

You cannot set attributes for DATA volumes.

Setting emulation for CKD volumes is not supported. If attempting to set attributes for multiple volumes of type FBA and CKD, a warning is displayed stating that the action will be applied only to FBA volumes.


Setting emulation is not supported on masked/mapped volumes.


To set volume emulation:



**Procedure**

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Select one of the volume type tabs.
4.
 



 Select a volume, click , and click **Set Volumes > Emulation**.
5. Select the Emulation type.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## FAST association operations


The follow FAST association operations are available:

- Associating storage groups with FAST policies (see [Associating storage groups with FAST policies](#) on page 169).
- Associating FAST policies with storage groups (see [Associating FAST policies with storage groups](#) on page 168).
- Re associating FAST polices and storage groups (see [Reassociating FAST polices and storage groups](#) on page 170).

## Removing DATA volumes

This procedure explains how to remove DATA volumes on storage systems running Enginuity version 5876.

**Procedure**

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click  to open its **Details** view.
4. Click the number next to **Number of Data Volumes**.
5. Select a data volume and click **Remove**.
6. Click **OK**.

## Mapping volume operations

The follow mapping volume operations are available:

- Mapping volumes (see [Mapping volumes](#) on page 192).
- Unmapping volumes (see [Unmapping volumes](#) on page 193).
- Mapping CKD volumes (see [Mapping CKD volumes](#) on page 340).

- Unmapping CKD volumes (see [Unmapping CKD volumes](#) on page 341).
- z/ OS map from the volume list view (see [z/OS map from the volume list view](#) on page 333).
- z/ OS unmap from the volume list view (see [z/OS unmap from the volume list view](#) on page 334).
- z/ OS map from the Volumes (Storage Groups) list view (see [z/OS map from the Volumes \(Storage Groups\) list view](#) on page 335 ).
- z/ OS unmap from the Volumes (Storage Groups) list view (see [z/OS unmap from the Volumes \(Storage Groups\) list view](#) on page 335).
- z/ OS map FBA volumes from the Volumes (Storage Groups) list view (see [z/OS map FBA volumes from the Volumes \(Storage Groups\) list view \(HYPERMAX OS 5977 or higher\)](#) on page 338).
- z/ OS map from the CU image list view (see [z/OS map from the CU image list view](#) on page 332).
- z/ OS unmap from the CU image list view (see [z/OS unmap from the CU image list view](#) on page 333).

## Rename operations

The follow rename operations are available:

- Rename disk groups (see [Renaming disk groups](#) on page 237).
- Rename storage tiers (see [Renaming tiers](#) on page 163).

## Provisioning storage

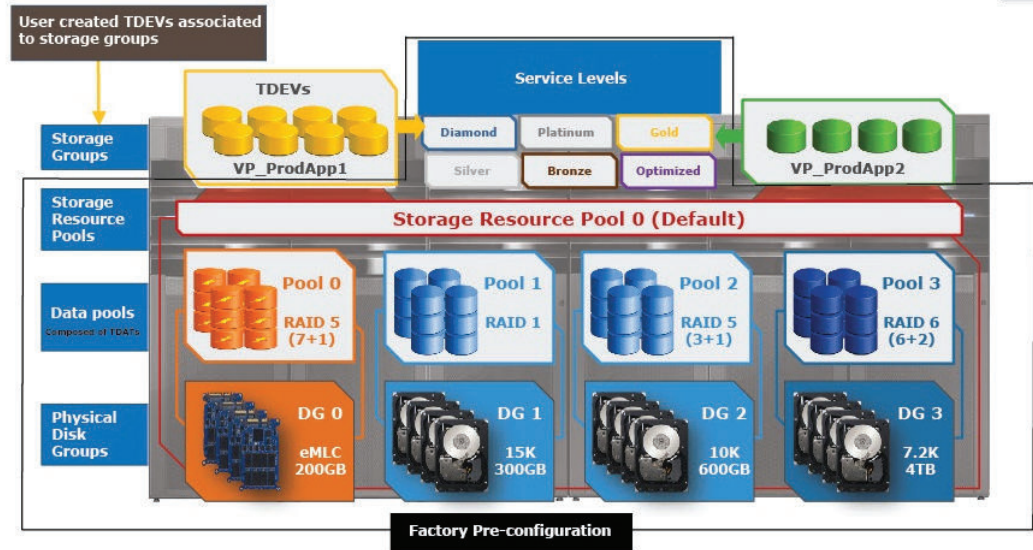
With the release of HYPERMAX OS 5977 and the next generation storage systems, Unisphere introduces support for service level provisioning. Service level provisioning simplifies storage management by automating many of the tasks associated with provisioning storage.

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

By default, storage 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. allows storage administrators to view all the SRPs configured on the system and the demand that storage groups are placing on them.

In addition, storage 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 , and workload to the storage group containing the application's LUNs.



Unisphere provides the following methods for provisioning 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, and so on).

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

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

Regardless of the method you choose, once you have completed the process 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 storage system must already be configured.

For instructions on provisioning storage systems running Enginuity 5876 or higher, refer to [Provisioning storage](#) on page 107.

To provision storage:

<ol style="list-style-type: none"> <li>1. <a href="#">Creating hosts</a> on page 292 Use the <b>Create Host</b> dialog box to group host initiators (HBAs).</li> <li>2. <a href="#">Using the Provision Storage wizard</a> on page 100 Use the <b>Provision Storage</b> wizard, which will step you through the process of creating the storage group, port group, and masking view.</li> </ol>	<ol style="list-style-type: none"> <li>1. <a href="#">Creating hosts</a> on page 292 Use the <b>Create Host</b> dialog box to group host initiators (HBAs).</li> <li>2. <a href="#">Creating volumes</a> on page 178 Create one or more volumes on the storage system.</li> <li>3. Use the <b>Create Storage Group</b> 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, and a workload.</li> <li>4. <a href="#">Creating port groups</a> on page 316 Group Fibre Channel and/or iSCSI front-end directors.</li> <li>5. <a href="#">Creating masking views</a> on page 307</li> </ol>
---	---

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

## Using the Provision Storage wizard

### Before you begin

- The storage system is running HYPERMAX OS 5977 or higher.
- 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 does not apply. Or, if you open the wizard from the Provisioning Templates view, the steps on selecting the Service Level and Workload Type does not apply. When opening the wizard from the Provisioning Templates view, please also note the following:

Based on the selected template, the appropriate fields (service level, workload type, size and number of volumes) will be filled in with values from the template. If the service level is not available on the default SRP on the selected storage system, it will default to the default service level (Diamond for AFA, Optimized for hybrid arrays).

When creating a storage group from the first page without adding it to a masking view, the storage group will be associated with the template but will be marked invalid (will not be included in the usage count for that template) until it is added to a masking view.

If the selected template has host IO limits defined based on the provisioning request the limits will be set

- Standalone SG: The limits will be set and can be modified.
- Cascaded SG: The limits will be set on each of the children but the parent will have no limit set.

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** will open the wizard on the Select Host/Host Group page because you are starting out with a storage group.

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.

The maximum number of storage groups allowed on a storage system running HYPERMAX OS 5977 is 16,384.



For HYPERMAX OS 5977 or higher, the maximum number of child storage groups allowed in a cascaded configuration is 64.

For instructions on provisioning storage systems running Enginuity 5876, refer to [Using the Provision Storage wizard](#) on page 108.

To use the Provisioning Storage wizard:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**.

3. Do one of the following:
  - Select the storage group and click **Create** to open the **Provision Storage** wizard.
  - Select the storage group and click **Provision Storage to Host** to open the **Provision Storage** wizard (go to step 8).
4. 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.
5. If required, select an Emulation type.
6. 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**.
7. Optional: Add one or more storage groups by hovering over the area to the right of the volume capacity and selecting .
8. Optional: Create a storage group with multiple volume sizes or edit the storage group by hovering over the area to the right of the volume capacity and selecting  (see [Editing storage group volume details](#) on page 149).
9. 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 critical, 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)	Optimized	Places the most active data on the highest performing storage and the least active on the most cost-effective storage.

For all-flash storage systems running HYPERMAX OS 5977, the only service level available is Diamond and it is selected by default.

10. Select the **Workload Type** to assign it.

---

#### Note

Workload type is not supported for CKD storage groups.

---

---

**Note**

Starting with Unisphere 9.0, workloads are not supported on PowerMaxOS 5978 and higher.

---

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).

11. Type the number of **Volumes** and select the **Capacity** of each.
- 

**Note**

The maximum volume size supported on a storage system running HYPERMAX OS 5977 is 64 TB.

It is possible to create an empty Storage Group with no volumes.

---

12. 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 the help page for the dialog box. When done, close the dialog box to return to the wizard.
13. Compression is enabled by default on All Flash systems running the HYPERMAX OS 5977 Q3 2016 Service Release or higher when you are creating a storage group or storage container. To disable it, uncheck the **Compression** check box. For more information, refer to [Understanding compression](#).
14. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
15. 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.
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**. 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**.

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: To receive alerts when the performance of the storage group changes, relative to its service level target, select **Enable Compliance Alerts**. For more information on Compliance Alerts, refer to [Creating service level compliance alerts policies](#) on page 63.
19. Optional: Click **Set Host I/O Limits**.  
This option is not displayed when you select an existing storage group and click **Provision**. The option is displayed when you click **Create**.
20. Optional: Determine if the storage system can handle the updated service level:
  - 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](#) on page 111.
  - Click **OK** to close the message.
  - If your updates are found to be unsuitable, modify the settings and run the check again until the suitability check passes.
21. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.



## Provisioning storage for mainframe

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

Service level provisioning eliminates the need for storage administrators to manually assign physical resources to their applications. Instead, storage administrators specify the service level and capacity required for the application and the system provisions the storage group appropriately.

You can provision CKD storage to a mainframe host using the Provision Storage wizard. For specific instructions about how to provision storage for mainframe, refer to [Using the Provision Storage wizard for mainframe](#) on page 104.

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

To provision storage for Open Systems, refer to [Using the Provision Storage wizard](#) on page 100.

### Mapping CKD devices to CU images

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

For more information about how to map CKD devices to CU images, see the following tasks:

- [z/OS map from the CU image list view](#) on page 332
- [z/OS map from the volume list view](#) on page 333

## Using the Provision Storage wizard for mainframe

### Before you begin

- The storage system must be running HYPERMAX OS 5977 Q1 2016, or higher, and have at least one FICON director configured.
- Depending on the type of configuration selected, not all of the steps listed below might be required.

To provision storage to mainframe:

### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.
3. In the Actions panel, click **Provision Storage**. The Provision Storage wizard for mainframe is displayed.
4. In the **Create Storage Group** page, 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.

If you want to create an empty storage group, proceed to the final step after typing the storage group name.



5. 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.

6. Select an **Emulation** type. Available values are **CKD-3390** and **CKD-3380**.

7. 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.

Available values are:

Service level	Performance level	Use case
Diamond	Ultra high	HPC, latency sensitive
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.

For all-flash storage systems, the only service level available is Diamond and it is selected by default.

8. Type the number of **Volumes** and select either a **Model** or **Volume Capacity**.

Selecting a **Model** type automatically updates the **Volume Capacity** value. Alternatively, you can type the **Volume Capacity**.

---

**Note**

The maximum CKD volume size supported is 1182006 cylinders or 935.66 GB. It is possible to create an empty Storage Group with no volumes.

---


9. (Optional) Configure volume options:

---

**Note**

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

---

- a. Hover the cursor on the service level and click  .

b. Edit the **Volume Identifier**.

The following options are available:

**None**

Do not set a volume identifier.

**Name Only**

All volumes will have the same name. Type the name in the **Name** field.

**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. Type the name in the **Name** field.

**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. Type the name in the **Name** field.

- c. To **Allocate capacity for each volume** you are adding to the storage group, select this option. You can use this option only for newly created volumes, not existing volumes.
  - d. 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.
  - e. Click **OK**.
10. (Optional) To add a child storage group, do one of the following:
- On all-flash storage systems, click **Add Storage Group**.
  - On all other storage systems click **Add Service Level**.

Specify a **Name**, **Service Level**, **Volumes**, and **Model/Volume Capacity**.

Repeat this step for each additional child storage group. The maximum number of child storage groups allowed is 64.

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:
- 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 920 and [Previewing jobs](#) on page 920
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
12. On the **CU Image** page, select whether to use a **New** or an **Existing** CU image, and then do the following depending on your selection:
- New:
    - a. Specify the following information for the new CU image:
      - **CU Image Number**
      - **SSID**
      - **Base Address**
    - b. Select a **Split** with which to associate the CU image.
  - Existing:
    - a. Select a CU image.

- b. To specify a new value for the base address, click **Set Base Address**. For more information about setting the base address, refer to [Setting the base address](#) on page 337.
13. Click **Next**.
14. On the **Review** page, review the summary information displayed.  
 If the storage system is registered for performance, you can subscribe for compliance alerts for the storage group and run a suitability check to ensure that the load being created is appropriate for the storage system.  
 To enable compliance alerts, select **Enable Compliance Alerts**.  
 To run a suitability check, click **Run Suitability Check**.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Provisioning storage

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


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

**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, and so on).

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

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

Regardless of the method you choose, once you have completed the process 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 storage system must already be configured.

To provision storage:

<ol style="list-style-type: none"> <li>1. Use the Create Host dialog box to group host initiators (HBAs).</li> <li>2. Use the Provision Storage wizard, which 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.</li> </ol>	<ol style="list-style-type: none"> <li>1. <a href="#">Creating hosts</a> on page 292 Use the Create Host dialog box to group host initiators (HBAs).</li> <li>2. <a href="#">Creating volumes</a> on page 178 Create one or more volumes on the storage system.</li> <li>3. Use the Create Storage Group wizard to create a storage group. If you want to add the volumes you created in step 2, be sure to set the wizard's Storage Group Type to Empty, and then complete <a href="#">Adding volumes to storage groups</a> on page 114.</li> <li>4. <a href="#">Creating port groups</a> on page 316 Group Fibre Channel and/or iSCSI front-end directors.</li> <li>5. <a href="#">Creating masking views</a> on page 307 Associate the host, storage group, and port group into a masking view.</li> <li>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.</li> </ol>
---	---

## Using the Provision Storage wizard

### Before you begin

The storage system is running Enginuity OS version 5876 and must already be configured and you must already have a host. For instructions on creating a host, refer to [Creating hosts](#) on page 292.

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.

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

The maximum number of storage groups allowed on a storage system running Enginuity 5876 is 8,192.

For Enginuity 5876 or higher, the maximum number of child storage groups allowed in a cascaded configuration is 32.

For users who want the ability to control every aspect of the provisioning process, refer to the Advanced procedure in [Using the Provision Storage wizard](#) on page 108.




For instructions on provisioning storage systems running HYPERMAX OS 5977, refer to [Using the Provision Storage wizard](#) on page 100.




To use the Provisioning Storage wizard:

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**.
3. Do one of the following:
  - Select the storage group and click **Create** to open the **Provision Storage** wizard.
  - Select the storage group and click **Provision Storage to Host** to open the **Provision Storage** wizard (go to step 8).
4. 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.

5. Select the **Storage Group Type**.
6. Do the following, depending on the storage group type:
  - Standard Storage Group:
    - Select the **Volume Type** to add to the storage group and click **NEXT**.
    - Do the following, depending on the volume type:
      - Virtual Volumes:
        - a. Select the **Emulation** type for the volumes to add to the storage group.
        - b. Optional: Select the **Thin Pools** containing the volumes to add to the storage group.
        - c. Type the number of volumes and enter volume capacity information.
        - d. Optional: To add more volumes, hover the cursor over the volume and click .
        - e. Optional: To remove a previously added volume, hover the cursor over it and click .
        - f. Optional: To edit a volume, hover the cursor over the volume and click edit (  ) (see [Editing storage group details](#) on page 150)
      - 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. Type the number of volumes and enter volume capacity information.
  - e. Optional: To add more volumes, hover the cursor over the volume and click .
  - f. Optional: To remove a previously added volume, hover the cursor over it and click .
  - g. Optional: To edit a volume, hover the cursor over the volume and click edit (  ) (see [Editing storage group details](#) on page 150)
- Empty Storage Group:  
Note: It is possible to create an empty Storage Group with no volumes.  
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
7. 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:  
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
8. Specify the host/host group to use by selecting an existing host/host group, or doing the following to create a new host or host group.
  - To create a new host, click **Create Host** (see [Creating hosts](#) on page 292).
  - To create a new host group, click **Create Host Group** (see [Creating host groups](#) on page 302).
9. Click **NEXT**.
10. 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), click **Include ports not visible to the host** slider. The following properties display:

- **Dir-Port**—Storage system director and port in the port group.
- **Identifier**—Identifier.
- **Initiators**—Number of initiators logged into the fabric.
- **PGs**—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**.

11. Optional: Edit the suggested name for the Masking View by highlighting it and typing a new name over it.
12. Optional: To set host I/O limits for the storage groups, click **Set Host I/O Limits**. For information about setting the limits, refer to [Setting host I/O limits](#) on page 132.

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.

13. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Suitability Check restrictions



The Suitability Check option is only available when:

- The storage system is running HYPERMAX OS 5977 or higher.
- The storage system is registered with the performance data processing option for statistics.
- The workloads have been processed.
- All the SGs involved have a service level and SRP set.
- The target SRP does not contain only external disk groups (like XtreamIO).
- The storage system is local.
- The SG is not in a making view (only for the local provisioning wizard). The SG should be in a masking view for the Modify SG case.

## Suitability Check

The Suitability Check option is only available when the storage system is running HYPERMAX OS 5977 or higher.

This message indicates whether the storage system can handle the updated service level. Results are indicated with either of the following:

-  Indicates suitable.
-  Indicates non-suitable.

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 (viewed by hovering the cursor over the bar) indicating the components expected availability on the target storage system after the change.

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.

## Creating storage groups

This procedure explains how to create storage groups on storage systems running HYPERMAX OS 5977 or later. 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](#) on page 100.

For instructions on creating storage groups on storage systems running Engenuity 5876, refer to [Using the Provision Storage wizard](#) on page 108.


Before you begin:

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

To create a storage group:

### Procedure

1. Select the storage system.
2. Select **Storage > Volumes**.
- 3.

Select one or more volumes, click  and select **Create SG**.

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 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 critical, 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.
None		

For all-flash storage systems running HYPERMAX OS 5977, the only service level available is Diamond and it is selected by default.

7. Refine the service level by selecting the **Workload Type** to assign it.

---

**Note**

Workload type is not supported for CKD storage groups.

---

**Note**

Starting with Unisphere 9.0, workloads are not supported on PowerMaxOS 5978 and higher.

---

Possible values for the **Workload Type** are:

- OLTP
- OLTP+REP
- DSS
- DSS+REP

The workload type does not apply when the service level is Optimized or None.

8. Click **OK** to create the storage group now, or click **Advanced Options** to continue setting the advanced options, as described in the remaining steps.
9. Compression is enabled by default on All Flash systems running the HYPERMAX OS 5977 Q3 2016 Service Release or higher when you are creating a storage group or storage container. To disable the feature, uncheck the **Compression** check box. For more information, refer to [Understanding compression](#)
10. Optional: Click the **Enable Mobility ID** checkbox to assign Mobility IDs to the volumes in the storage group. If you leave the checkbox unchecked, Compatibility IDs will be assigned to the volumes instead.
11. Optional: Select **Allocate Full Volume capacity**.
12. Optional: Click **Persist preallocated capacity through reclaim or copy** checkbox.

13. 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.

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Adding volumes to storage groups

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


Before you begin:

A storage group can contain up to 4,096 volumes.

A volume can belong to more than one storage group.

To add volumes to storage groups:

### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
3. Select the storage group and click  .
4. Click the number next to **Volumes**.
5. Click **Add Volumes to SG** to open the **Add Volumes to Storage Group** wizard.
6. Locate the volumes by selecting or typing values for any number of the following criteria:
  - **Capacity equal to**—Filters the list for volumes with a specific capacity and capacity type.
  - **Volume ID**—Filters the list for a volume with specific ID.
  - **Volume Identifier Name**—Filters the list for the specified volume name.
  - **Volume configuration**—Filters the list for the specified volume configuration.
  - **Emulation**—Filters the list for the specified volume emulation.
  - **Exclude Volumes in use**—Tick the checkbox to filter the list to exclude volumes in use.
7. Click **NEXT** to run the query.
 

Results are displayed on the next page in the wizard.
8. Select the volumes and click **OK**.

## Copying volumes between storage groups



This procedure explains how to copy volumes between storage groups.

Before you begin:

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

To copy volumes between storage groups:

#### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
3. Select the storage group and click  .
4. Click **Volumes**.
5. Select one or more volumes click  , and click **Copy Volumes To SG** to open the **Copy Volumes to Storage Group** dialog box.
6. Select the **Target Storage Group Name**.
7. Click **OK**.

## Moving volumes between storage groups



This procedure explains how to move volumes between storage groups.

Before you begin:

- Storage groups require Enginuity 5876 or HYPERMAX OS 5977 or later.
- 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:

#### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
3. Select the storage group and click  .
4. Click **Volumes**.
5. Select one or more volumes, click  , and click **Move Volumes to SG** to open the **Move Volumes to Storage Group** dialog box.
6. Select the **Target Storage Group Name**.
7. 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**.

8. Click **OK**.

## Removing volumes from storage groups


This procedure explains how to remove volumes from storage groups.

Before you begin:

Storage groups require Enginuity 5876 or HYPERMAX OS 5977 or higher.

To remove volumes from storage groups:

### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
3. Select the storage group and click  .
4. Click **Volumes**.
5. Select one or more volumes and click **Remove Volumes** to open the **Remove Volume** dialog box.
6. To unbind the volumes, select **Unbind** or **Unmap**, depending on the storage operating environment.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Storage Group operations

The following storage group operations are available:

- Expanding Storage Group (see [Expanding storage groups](#) on page 116).
- Modifying Storage Group (see [Modifying storage groups](#) on page 119 (5977 or greater)).

## Expanding storage groups

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 OS 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 196.
- Empty SGs are not displayed while creating a cascaded SG.

To expand a storage group:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**.
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.
  - **Copy Volume**—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 Pools** containing the volumes to add to the storage group.
    - c. Type the number of volumes and enter volume capacity information.
    - d. Optional: To add more volume sizes, hover the cursor over the volume and click .
    - e. Optional: To remove a previously added volume, hover the cursor over it and click .
    - f. Optional: To edit a volume, hover the cursor over the volume and click edit (  ) (see [Editing storage group details](#) on page 150)
  - 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. Type the number of volumes and enter volume capacity information.
    - e. Optional: To add more volume sizes, hover the cursor over the volume and click .
    - f. Optional: To remove a previously added volume, hover the cursor over it and click .
    - g. Optional: To edit a volume, hover the cursor over the volume and click edit (  ) (see [Editing storage group details](#) on page 150)
  - Copy Volume:
    - 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 entering volume capacity information.
    - e. Optional: Hover the cursor over the volume and click edit (  ) (see [Editing storage group details](#) on page 150)
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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Expanding ProtectPoint storage groups

### 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 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 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.
- CKD devices are not supported by ProtectPoint.

This procedure explains how to increase the amount of storage in a storage group protected by ProtectPoint.

To expand protected storage groups:

### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
- 3.

Select the storage group, click , and click **Expand ProtectPoint**.

Opens the **Expand ProtectPoint** wizard.

4. Select the Point In Time Copy to expand and click **Next**.
5. Select the external LUNs to add to the backup storage group and click **Add to Group**. Select the same number of external LUNs as the number of volumes added to the production storage group.
6. Click **Next** and select the Restore Storage Group.
7. Select the external LUNs to add to the restore storage group and click **Add to Group**. Select the same number of external LUNs as the number of volumes added to the production storage group.
8. Click **Next** and verify your selections. To change any of them, click **Back**. Some changes may require additional configuration changes.
9. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now**.

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.

## Modifying storage groups

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

Before you begin:

- You must be an Administrator or StorageAdmin.
- The maximum number of storage groups allowed on a storage system is 16,384.
- 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:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**
3. Select the storage group and click **Modify** to open the **Modify Storage Group** dialog box.
4. 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 compliance alerts policies](#) on page 65.
  - 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 levels 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

Service level	Performance type	Use case
Platinum	Very high	Mission critical, 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.

For all-flash storage systems running HYPERMAX OS 5977, the only service level available is Diamond and it is selected by default.

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

---

#### Note

Starting with Unisphere 9.0, workloads are not supported on PowerMaxOS 5978 and higher.

---

- 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.

---

#### Note

In mixed FBA/CKD All Flash systems, volume capacity defaults to **GB** for FBA Storage Groups and **Cyl** for CKD Storage Groups.

---

- 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 storage system running HYPERMAX OS 5977 is 64 TB.

All Flash systems running the HYPERMAX OS 5977 Q2 2017 Service Release or higher supports a maximum CKD device size of up to 1,182,006 cylinders.

---

- g. SRDF storage group volume capacity can be expanded using the controls. In the case of SRDF Storage Groups, you need to specify a SRDF group



- number so that the dialog allowing you to remote volumes can also be displayed (see [Expanding remote volumes](#) on page 511).
- h. Optional: Add one or more storage groups by hovering over the area to the right of the volume capacity and selecting  .
  - i. Optional: Create a storage group with multiple volume sizes or edit the storage group by hovering over the area to the right of the volume capacity and selecting  (see [Editing storage group volume details](#) on page 149).
  - j. Optional: to add a child storage group, do one of the following:
    - On all-flash storage systems, click **Add Storage Group**.
    - On all other storage systems, click **Add Service Level**.
    - Modify any of the service level parameters, as described earlier in this procedure.
5. Compression is enabled by default on All Flash systems running the HYPERMAX OS 5977 Q3 2016 Service Release or higher when you are creating a storage group or storage container. To disable the feature, uncheck the **Enable Compression** check box. In a cascaded setup, changes will be passed to each of the child storage groups. For more information on compression, refer to [Understanding compression](#)
  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](#) on page 111.
    - 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. 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Renaming storage groups

This procedure explains how to rename storage groups.


Before you begin:


- Storage group names must be unique from other storage groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores ( \_ ), and (-) are allowed. Storage group names are case-insensitive.
- Storage groups require Enginuity 5876, or HYPERMAX OS 5977 or later.

To rename a storage group:

**Procedure**

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
3.
 



 Select the storage group, click , and click **Rename**.
4. Type the new name.
5. Click **OK**.

**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 387. This is the default method for storage systems running HYPERMAX OS 5977 or higher.
- **TimeFinder/Clone**—For instructions, refer to [Protecting storage groups using TimeFinder/Clone](#) on page 122. This is the default method for storage systems running Enginuity 5876.
- **ProtectPoint**—For instructions, refer to [Protecting storage groups using ProtectPoint](#) on page 124. This method is only available on storage systems running HYPERMAX OS 5977 or later.
- **RecoverPoint**—For instructions, refer to [Protecting storage groups using RecoverPoint](#) on page 125. This method is only available for storage systems running Enginuity 5876.
- **SRDF**—For instructions, refer to [Protecting storage groups using SRDF](#) on page 126. 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 127. This method is only available for storage systems running HYPERMAX 5977 or higher.

**Protecting storage groups using TimeFinder/Clone**

Before you begin:

- This feature requires the Enginuity 5876.163.105 or later. This feature does not apply to storage systems running HYPERMAX OS 5977 or later.
- The storage group must contain only thin volumes (except gatekeepers under 10 MB) and they must all be of the same type (either BCV or standard thin volumes (TDEVs)). This restriction also applies to cascaded storage groups, that is, all volumes in the parent and child storage groups must be thin and of the same type.
- The SYMAP\_ALLOW\_DEV\_INT\_MULTI\_GRP option must be enabled. For instructions on enabling the option, refer to "Editing the Options file" in the *Solutions Enabler Installation Guide*.
- Meta volumes are not supported.

To protect storage groups using TimeFinder/Clone:

**Procedure**

1. Select the storage system.
2. Select **STORAGE > Storage Groups**.

3. Select the storage group and click **Protect**.
4. If not already selected, select **Point In Time 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 it's 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. 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. Select the mode in which to create the clone session **Clone Copy Mode**.  
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.
11. Select the type of volumes to use as the targets (**Clone Targets**).
12. Click **NEXT**.
13. Verify your selections, and then 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Protecting storage groups using ProtectPoint

### 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.
- CKD devices are not supported by ProtectPoint.

To protect storage groups using ProtectPoint:

### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
3. Select the storage group and click **Protect**.
4. Select **Backup Using ProtectPoint**.
5. Click **NEXT**.
6. Click **OK**.
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 920 and [Previewing jobs](#) on page 920.
  - 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 RecoverPoint

### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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 StorageAdmin.
- 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 4.1 is setup and operational. For each cluster in the setup, gatekeepers and repository volumes must be configured in their relevant masking view. uses a default journal masking view naming convention.
- 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.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**.
3. Select the storage group and click **Protect**.
4. On the **Select Technology** page, select **Remote Replication using RecoverPoint**.
5. Click **NEXT**.
6. On the **Configure RecoverPoint** 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.
  - **Data Initiator Group**—Data initiator group.
  - **Journal Thin Pool**—Journal thin pool.
  - **Journal Port Group**—Journal port group.
  - **Data Initiator Group**—Journal initiator group.
7. Click **NEXT**.
8. On the **Add Copies** page, specify the following information:
  - **RecoverPoint Cluster**—RecoverPoint cluster.

- **Copy Name**—Name of the RecoverPoint copy.
  - **Mode**—Specify whether the mode is Synchronous or Asynchronous.
  - **Array**—Storage system.
  - **Target Storage Group**—Specify whether the RecoverPoint copy targets a new storage group or an existing group.
  - **Copy Storage Group**—Name of storage group to be copied.
  - **Data Thin Pool**—Name of data thin pool.
  - **Data Port Group**—Name of data port group.
  - **Journal Thin Pool**—Name of journal thin pool.
  - **Journal Port Group**—Name of journal port group.
9. Click **Add Copy**.  
Lists the copy in the **Copy Summary** table.
  10. Click **NEXT**.
  11. On the **FINISH** page, verify your selections. To change any of them, click **BACK**. 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, then click **Run Now** to perform the operation now.

## Protecting storage groups using SRDF

This procedure explains how to protect storage groups using SRDF.

Before you begin:

- You must have StorageAdmin permission.
- Connectivity to remote storage system must be available.
- All storage systems involved must be discoverable and manageable from the console.
- The SRDF wizard in Unisphere 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.
- The SRDF wizard in Unisphere 8.2 and higher releases supports the creation of SRDF protection for CKD Storage Groups.
- Set the default number of ports to use with SRDF. To set this number, refer to [Managing data protection preferences](#) on page 85.

To protect storage groups using SRDF:

### Procedure

1. Select the storage system.

2. Under **STORAGE**, select **Storage Groups**.
3. Select the storage group and click **Protect**.
4. Select **Remote Replication Using SRDF**.
5. Click **NEXT**.
6. Select the Remote Array ID. To update the list of remote systems, click **Scan**.
7. Select the **Replication Mode**. For more information, refer to [SRDF session modes](#) on page 442.
8. Select **Automatic** to automatically select a SRDF group or **Manual** to select a SRDF group from a list.
9. Optional: To not start pair mirroring, clear the **Establish Pairs** option.
10. Do the following, depending on the storage operating environment (target system):

For HYPERMAX OS 5977 or later:

Optional: Change the **Remote Storage Group Name**, and optionally select a **Remote Service Level**. Changing the name will also change the target volume storage group name.

For Enginuity 5876:

- a. Optional: Change the **Remote Storage Group Name**.
  - b. 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 it is 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.
  - c. Optional: Select the **Remote FAST Policy**. This is the FAST policy associated with the remote storage group.
  - d. 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.
11. For HYPERMAX OS 5977 or later, click **Create Device Group** check box and select the **Device Group Name** that will hold the target volumes.
  12. Click **NEXT**.
  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 920 and [Previewing jobs](#) on page 920.
    - 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 later.
- You must have StorageAdmin permission.
- Connectivity to remote storage system must be available.
- All storage systems involved must be discoverable and manageable from the console.
- CKD devices are not supported by SRDF/Metro.

You are not allowed to set RDF devices in the non-Metro RDF mirror to operate in Synchronous mode.

For systems running PowerMaxOS 5978 or higher, the create pair operation is blocked if the device ID types of each individual SRDF device pair are not the same (both Compatibility ID or both Mobility ID) on both sides. Device type ID conversion from a Compatibility ID to a Mobility ID is not allowed on a device once it is part of an SRDF/Metro session. Candidate IDs are restricted to those running PowerMaxOS 5978 or higher if the source storage group has devices with Mobility ID in them.

Protecting a storage group using SRDF/Metro from the protection wizard is allowed when one or more of the devices in the storage group have the GCM flag set.

To protect storage groups using SRDF Metro:

#### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
3. Select the storage group and click **Protect**.  
The **Select Protection Type** page displays.
4. Select **High Availability using SRDF/Metro**.
5. Click **NEXT**.
6. Select the **Remote Array ID**. To update the list of remote arrays, click **Scan**.
7. Optional: To stop the initiation of pair mirroring, clear the **Establish Pairs** option.
8. If **Establish Pairs** is checked, choose Protected by **Witness** or **Bias**.

If Witness is unavailable on the local or remote array, the option is disabled and Bias is selected by default. If available, Witness is selected by default.

For storage systems running HYPERMAX OS 5977 Q3 2016 or higher, when the **Witness** radio button is selected, the Witness Candidate (Remote Array) field displays a list of physical and Virtual witnesses instances which are enabled. Disabled virtual witness instances are not displayed.

9. Optional: Change the **Remote Storage Group Name**, and optionally select a **Remote Service Level**.  
Changing the name will also change the target volume storage group name.
10. Optional: To disable compression, clear the **Compression** option.
11. Click **NEXT**.
12. Verify your selections. To change any of them, click **BACK**.

Note that some changes may require you to make additional changes to your configuration.

13. 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 920 and [Previewing jobs](#) on page 920.
- 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.


Before you begin:

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

To convert storage groups:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**
- 3.

Select the storage group, click , and click **SG Maintenance > Convert to Cascaded**.

4. 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.

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Changing Storage Resource Pools for storage groups

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

In eNAS environments, you can also perform this operation from the File Storage Groups page (**System > System Dashboard > File Dashboard > File Storage Groups**).

Before you begin:


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

To change the Storage Resource Pool for storage groups:

### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.

3.

Select the storage group, click , and select **Change SRP** to open the **Change SRP** dialog box.

4. Select the new SRP.

5. (Optional) Change the Service Level for the SG. Service levels 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 critical, 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.

For all-flash storage systems, the only service level available is Diamond and it is selected by default.

6. (Optional) Refine the service level by selecting the **Workload Type** to assign it. (This step is not applicable for storage systems running PowerMaxOS 5978.)
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.


## Adding or removing cascaded storage groups

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

Before you begin:

To add or remove cascaded storage groups:

### Procedure

- Select the storage system.
- Select **Storage > Storage Groups** to open the **Storage Group** list view.
- Select the parent storage group and click  to open its **Details** view.
- Click the number next to **Storage Groups** to open the child **Storage Groups** list view.

5. Do the following, depending on whether you are adding or removing storage groups:
  - Adding storage groups:
    - a. Click **Add**.
    - b. Select one or more storage groups.
    - c. Do one of the following:
      - a. 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 920 and [Previewing jobs](#) on page 920.
      - b. Expand **Add to Job List**, and click **Run Now** to perform the operation now.
  - Removing storage groups:
    - a. Select one or more storage groups and click **Remove**.
    - b. Click **OK**.

## Renaming storage groups

This procedure explains how to rename storage groups.


Before you begin:

- Storage group names must be unique from other storage groups on the storage system and cannot exceed 64 characters. Only alphanumeric characters, underscores ( \_ ), and (-) are allowed. Storage group names are case-insensitive.
- Storage groups require Enginuity 5876, or HYPERMAX OS 5977 or later.

To rename a storage group:

### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
- 3.

Select the storage group, click , and click **Rename**.

4. Type the new name.
5. Click **OK**.

## Deleting storage groups

This procedure explains how to delete storage groups.



Before you begin:

- Storage groups require Enginuity 5876 or HYPERMAX OS 5977 or later.
- 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 compliance alerts policies](#) on page 65.

To delete a storage group:

#### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
3.  Select the storage group, click , and select **Delete**.
4. Click **OK**.

## 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 later, or HYPERMAX OS 5977 or later.
- 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 later, 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:

#### Procedure

1. Select a storage system.
2. Under **STORAGE**, select **Storage Groups**.

3. Select the storage group and select **Set Host I/O Limits** to open the **Set Host I/O Limits** dialog box.
4. Select and type values for one or both of the following:
  - **MB/Sec**—Maximum bandwidth (in MB per second). Valid values range from 1 MB/sec to 100,000 MB/sec.
  - **IO/Sec**—Maximum IOPs (in I/Os per second). Valid values range from 100 IO/Sec to 2,000,000 IO/sec, in 100 increments.
5. To configure a dynamic distribution of host I/O limits, set **Dynamic Distribution** to one of the following; otherwise, leave this field set to Never (default). This feature requires Enginuity 5876.163.105 or later.
  - **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.
6. Click **OK**

## Splitting storage groups

This procedure explains how to split cascaded storage groups on storage systems.


Unisphere supports the splitting of storage groups in two different ways:

- During the split operation, a specified child storage group is removed from the parent storage group. A new masking view is created on this child storage group with the same initiator groups and port groups of the parent storage group masking view.
- During the split operation, a new storage group with the user specified name will be created. The user specified devices from the source standalone storage group are moved to the newly created storage group and a new masking view is created on the new storage group using the same initiator groups and port groups of the source standalone storage group masking view.

To split a storage group:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**
- 3.

Select the storage group, click , and click **SG Maintenance > Split From**.

4. Do one of the following:
  - When splitting a child storage group from its parent masking view and moving it to a standalone masking view, select a child storage group and specify a new masking view name.
  - When splitting a standalone storage group into two storage groups each with a masking view, specify a new storage group name, specify a new masking view name, and select the volumes to be added to the new storage group.

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Merging storage groups


This procedure explains how to merge storage groups in order to create a cascaded storage group.


Unisphere support merging the masking views of a source standalone storage group and a target storage group which has a common initiator group and port group. The target storage group may be a parent storage group or a standalone storage group. In the case of the target being a parent storage group, during the merge operation, the source standalone SG is added to the target parent storage group and uses the parent storage group masking view. In the case of the target storage group being a standalone storage group, all devices in the source standalone storage group are moved to the target storage group. The source standalone storage group and its masking view are deleted.

To split a storage group:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**
3.
 



 Select the storage group, click , and click **SG Maintenance > Merge Into**.
4. Select the target storage group.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Managing VP compression on thin volumes in 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 storage systems running HYPERMAX OS 5977 or later.

To manage VP compression on storage groups:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** panel.
3.
 



 Select a volume, click , and click **VP Compression**.

4. Select one of the following compression operations:
  - **CompressStart**—Starts compressing the thin volumes in the storage group.
  - **CompressStop**—Stops compressing the thin volumes in the storage group.
  - **UnompressStart**—Starts uncompressing the thin volumes in the storage group.
  - **UncompressStart**—Stops uncompressing the thin volumes in the storage group.
5. Click **OK**.

## Viewing storage groups






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.

For information on viewing cascaded storage groups, see [Viewing cascaded storage groups](#) on page 141.

### Procedure

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

The following properties display:

- **Name** — Name of the storage group.
- **Compliance** — How well the storage group is complying with its service level, if applicable. Possible values are:
  -  Critical—Storage group is performing well below service level targets.
  -  Marginal—Storage group is performing below service level target.
  -  Stable—Storage group is performing within the service level target.
  -  Storage group has no assigned service level.
  -  Compliance information is being collected.
- **SRP** — Name of SRP that the storage group belongs to, if any.
- **Service Level**—Name of the service level associated with the storage group. If there is no service level associated with the group, then file displays N/A.
- **Capacity (GB)**—Total capacity of the storage group in GB.
- **Emulation**—Emulation associated with the storage group. .

The following controls are available:

-  — [Viewing storage group details](#) on page 140.

- **Create** — [Using the Provision Storage wizard](#) on page 100.
- **Modify**—[Modifying storage groups](#) on page 119.
- **Provision**—[Using the Provision Storage wizard](#) on page 100.
- **Protect** —[Protecting storage groups](#) on page 122.
- **Set Host I/O Limits**—[Setting host I/O limits](#) on page 132.
- **Set Volumes > Set Volume Status**—[Setting volume status](#) on page 194.
- **Set Volumes > Replication QoS**—[Setting copy pace \(QoS\) for storage groups](#) on page 197.
- **Migrate**—[Creating a non-disruptive migration \(NDM\) session](#) on page 502
- **Allocate/Free/Reclaim > Start**—[Managing thin pool allocations](#) on page 244
- **Allocate/Free/Reclaim > Stop**—[Managing thin pool allocations](#) on page 244
- **SG Maintenance > Convert to Cascaded**—[Converting storage groups to cascaded](#) on page 129.
- **SG Maintenance > Split From**—[Splitting storage groups](#) on page 133.
- **SG Maintenance > Merge Into**—[Merging storage groups](#) on page 134.
- **SG Maintenance > Remove**—[Adding or removing cascaded storage groups](#) on page 130.
- **Change SRP**—[Changing Storage Resource Pools for storage groups](#) on page 129.
- **Delete**—[Deleting storage groups](#) on page 131.
- **Expand ProtectPoint**—[Managing thin pool allocations](#) on page 244.

## Viewing storage groups

This procedure explains how to view storage groups on a storage system running Enginuity 5876. There are multiple ways to view the same information. Depending on the method you use, some of the properties and controls may not apply.

To view storage groups associated with a FAST policy, see [Viewing storage group for FAST policies](#) on page 171.

For information on viewing cascaded storage groups, see [Viewing cascaded storage groups](#) on page 141.

### Procedure

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 storage system.

The following properties display:

- **Name** — Name of the storage group.
- **FAST Policy**—Policy associated with the storage group.
- **Capacity (GB)**—Total capacity of the storage group in GB.
- **Emulation**—Emulation type.
- **Masking Views** — Number of masking views associated with the storage group.



The following controls are available:

-  — [Viewing storage group details](#) on page 140.
- **Create**—[Using the Provision Storage wizard](#) on page 108.
- **Expand**—[Expanding storage groups](#) on page 116 .
- **Provision Storage to Host**—[Using the Provision Storage wizard](#) on page 108.
- **Protect**—[Protecting storage groups](#) on page 122.
- **Set Host I/O Limits**—[Setting host I/O limits](#) on page 132.
- **FAST > Associate** —[Associating FAST policies with storage groups](#) on page 168.
- **FAST > Disassociate**—[Disassociating FAST policies and storage groups](#) on page 170.
- **FAST > Reassociate**—[Reassociating FAST polices and storage groups](#) on page 170.
- **FAST > Pin**—[Pinning and unpinning volumes](#) on page 173.
- **FAST > Unpin**—[Pinning and unpinning volumes](#) on page 173 .
- **FAST > Bind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257.
- **FAST > Unbind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257.
- **FAST > Rebind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257.
- **Migrate**—[Creating a non-disruptive migration \(NDM\) session](#) on page 502
- **Allocate/Free/Reclaim > Start**—[Managing thin pool allocations](#) on page 244
- **Allocate/Free/Reclaim > Stop**— [Managing thin pool allocations](#) on page 244
- **SG Maintenance > Split From**—[Splitting storage groups](#) on page 133.
- **SG Maintenance > Merge Into**—[Merging storage groups](#) on page 134.
- **RecoverPoint > Tag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472.
- **RecoverPoint > Untag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472.
- **Delete**—[Deleting storage groups](#) on page 131.
- **Rename**—[Renaming storage groups](#) on page 121.
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945.
- **VP Compression**—[Managing VP compression on thin volumes in storage groups](#) on page 134.
- **Replication QOS**—[QOS for replication](#) on page 197.
- **Assign Symmetrix Priority**—[Assigning array priority to individual volumes](#) on page 189
- **VLUN Migration**—[Migrating regular storage group volumes](#) on page 261.
- **Set Optimized Read Miss**—[Setting optimized read miss](#) on page 193.


## Storage Group details

- Viewing storage group details on storage systems running HYPERMAX OS 5977 or later (see [Viewing storage group details](#) on page 138).
- Viewing storage group details on storage systems running Enginuity OS 5876 (see [Viewing storage group details](#) on page 140).

## Viewing storage group details

This procedure explains how to view configuration details for storage groups on storage systems running HYPERMAX OS 5977 or later. To view storage groups on a storage system running Enginuity OS 5876, refer to [Viewing storage group details](#) on page 140. 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.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click .

The following properties display:

- **SRP**—Name of SRP that the storage group belongs to, if any.
- **Compliance**—How well the storage group is complying with its service level, if applicable.
- **Service Level**—Service level associated with the storage group. If there is no service level associated with the group, then this field displays N/A.
- **Volumes**—Number of volumes in the storage group.
- **Child Storage Groups**—Number of child storage groups.
- **Masking Views**—Number of masking views associated with the storage group.
- **SnapVX Snapshots**—Number of SnapVX snapshots associated with the storage group.
- **SRDF**—SRDF.
- **Symmetrix ID**—Name of the storage group.
- **Capacity (GB)**—Total capacity of the storage group in GB.
- **VP Saved**—The percentage of space saved on the storage group.
- **Compression**— If compression is enabled on this storage group a tick will appear. If it's disabled a horizontal dash will appear.
- **Compression Ratio**— Current compression ratio for 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 132.
- **Host I/O Limit (MB/Sec)**— Maximum bandwidth (in MB per second). Valid values range from 1 MB/sec to 100,000 MB/sec.

- **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.
- **Emulation**—Emulation type.
- **Workload Type**—Workload type.
- **Dynamic Distribution**— 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.
- **Is Child**—Indicates whether the storage group is or is not a child storage group.
- **Parent Storage Group(s)**—Number of storage groups of which this storage group is a child. This field only displays for child storage groups.
- **RecoverPoint**—Indicates RecoverPoint usage.

Links are also provided to views for 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 the number next to **Volumes** will open a view listing the volumes contained in the storage group.

#### 4. Click **VIEW ALL DETAILS**.

A view with two tabs, **Details** and **Volumes** is displayed. Clicking the **Volumes** tab displays a view of the volumes in the storage group (see [Viewing volumes in storage groups](#) on page 142). Clicking the **Details** displays a view with two panels, a **Properties** panel and a **Capacity** panel.

The **Properties** panel displays the following:


- **Symmetrix ID**—Name of the storage group.
- **Compliance**—How well the storage group is complying with its service level, if applicable.
- **Service Level**—Service level associated with the storage group. If there is no service level associated with the group, then this field displays N/A.
- **Workload Type**—Type of the workload associated with the storage group.
- **SRP**—Storage resource pool (SRP) containing the storage group.
- **Masking Views**—Number of masking views associated with the storage group.
- **Emulation**—Emulation type.
- **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 132.
- **SnapVX Snapshots**—Number of SnapVX snapshots associated with the storage group.
- **SRDF**—SRDF.
- **Is Child**—Indicates that the storage group is or is not a child storage group.
- **Child Storage Groups**—Number of child storage groups.
- **RecoverPoint**—Indicates RecoverPoint usage.

The **Capacity** panel displays the following:

- **Capacity (GB)**—Total capacity of the storage group in GB.

- **Volumes**—Number of volumes in the storage group.
- **Allocated Capacity**—Number of volumes in the storage group.
- **VP Saved**— The percentage of space saved on the storage group.
- **Compression**— If compression is enabled on this storage group a tick will appear. If it's disabled a horizontal dash will appear.
- **Compression Ratio**— Current compression ratio for the storage group.

The following controls are available:


-  —[Storage Group operations](#) on page 116.
- **Set Host I/O Limits**—[Setting host I/O limits](#) on page 132.

## Results

### Viewing storage group details

This procedure explains how to view storage groups on a storage system running Enginuity OS 5876. There are multiple ways to view the same information. Depending on the method you use, some of the properties and controls may not apply.

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Groups** list view.
3. Select the storage group and click .

The following properties display:

- **Symmetrix ID**—Identity of the storage system.
- **FAST Policy**—Policy associated with the storage group.
- **Capacity (GB)**—Capacity of the storage group in GB.
- **Volumes**—Number of volumes in the storage group.
- **Child Storage Groups**—Number of child storage groups.
- **Masking Views**—Number of masking views associated with the storage group.
- **SRDF**—Number of SRDFs associated with the storage group.
- **Emulation** —Emulation type.
- **VP Saved**—The percentage of space saved on 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 132.
- **Host I/O Limit (MB/Sec)**—Maximum bandwidth (in MB per second). Valid values range from 1 MB/sec to 100,000 MB/sec.
- **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.
- **Dynamic Distribution**— 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.

- **Is Child**—Indicates whether the storage group is or is not a child storage group.
- **Parent Storage Group(s)**—Number of storage groups of which this storage group is a child. This field only displays for child storage groups.

Links are also provided to views for 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 the number next to **Volumes** will open a view listing the volumes contained in the storage group.

#### 4. Click **VIEW ALL DETAILS**.

A view with two tabs, **Details** and **Volumes** is displayed. Clicking the **Volumes** tab displays a view of the volumes in the storage group (see [Viewing volumes in storage groups](#) on page 142). Clicking the **Details** displays a view with two panels, a **Properties** panel and a **Capacity** panel.


The **Properties** panel displays the following:

- **Symmetrix ID**—Name of the storage group.
- **Masking Views**—Number of masking views associated with the storage group.
- **Emulation**—Emulation type.
- **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 132.
- **SRDF**—SRDF.
- **Is Child**—Indicates that the storage group is or is not a child storage group.
- **Child Storage Groups**—Number of child storage groups.
- **RecoverPoint**—Indicates RecoverPoint usage.

The **Capacity** panel displays the following:

- **Capacity (GB)**—Total capacity of the storage group in GB.
- **Volumes**—Number of volumes in the storage group.
- **Allocated Capacity**—Number of volumes in the storage group.
- **VP Saved**—The percentage of space saved on the storage group.

The following controls are available:

-  —[Storage Group operations](#) on page 116.
- **Set Host I/O Limits**—[Setting host I/O limits](#) on page 132.

## Results

### Viewing cascaded storage groups

#### Procedure

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

3. Select the storage group and click ⓘ to open its **Details** view.
4. Click the number next to **Storage Groups** to open the **Child Storage Groups** list view.
5. Optional: Use the **Child Storage Groups** list view to view and manage cascaded storage groups.
6. The following properties (depending on the storage operating environment) display:
  - **Name**—Name of the storage group.
  - **Compliance**—Indicates Compliance status.
  - **SRP**—SRP associated with the storage group.
  - **Service Level**—Service level associated with the storage group.
  - **Capacity (GB)**—Total capacity of the storage group in GB.
  - **Emulation**—Emulation type..
  - **Masking Views** —Number of masking views associated with the storage group.
7. The following controls are available:
  - ⓘ —[Viewing storage group details](#) on page 140
  - **Add** —[Adding or removing cascaded storage groups](#) on page 130
  - **Remove**—[Adding or removing cascaded storage groups](#) on page 130


## Viewing volumes in storage groups

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Group** list view.
3. Select the storage group and click ⓘ to open its **Details** view.
4. Click the number next to **Volumes** to open the **Volumes** list view.  
Use the **Volumes** list view to view and manage the volumes in a storage group.
5. The following properties display:
  - **Volume**—Assigned volume name.
  - **Type**—Type of volume.
  - **Allocated %**—Percentage allocated.
  - **Capacity (GB)**—Volume capacity in Gigabytes.
  - **Emulation**—Emulation type for the volume.
  - **Status** —Volume status.
  - **Pinned**—Whether the volume is pinned. Pinning volumes prevents any automated process such as FAST or Optimizer from moving them.

The following controls are available, depending on the storage operating environment:

- **Create** —[Creating volumes](#) on page 178.
- **Add Volumes to SG**—[Adding volumes to storage groups](#) on page 114.
- **Remove Volumes**—[Removing volumes from storage groups](#) on page 116.
- **Expand**—[Expanding existing volumes](#) on page 191
- **Copy Volumes to SG**—[Copying volumes between storage groups](#) on page 114.
- **Move Volumes to SG**—[Moving volumes between storage groups](#) on page 115.
- **Set Volumes > Emulation**—[Setting volume emulation](#) on page 96.
- **Set Volumes > Set Volume Attributes**—[Setting volume attributes](#) on page 195.
- **Set Volumes > Set Volume Identifiers**—[Setting volume identifiers](#) on page 196.
- **Set Volumes > Set Volume Status**—[Setting volume status](#) on page 194.
- **Set Volumes > Replication QoS**—[QOS for replication](#) on page 197.
- **Allocate/Free/Reclaim > Start**—[Managing thin pool allocations](#) on page 244 .
- **Allocate/Free/Reclaim > Stop**—[Managing thin pool allocations](#) on page 244 .
- **Configuration > Change Volume Configuration**—[Changing volume configuration](#) on page 190.
- **Configuration > Map**—[Mapping volumes](#) on page 192.
- **Configuration > Unmap**—[Unmapping volumes](#) on page 193.
- **Configuration > z/OS Map**—[z/OS map from the Volumes \(Storage Groups\) list view](#) on page 335 and [z/OS map FBA volumes from the Volumes \(Storage Groups\) list view \(HYPERMAX OS 5977 or higher\)](#) on page 338.
- **Configuration > z/OS Unmap**—[z/OS unmap from the Volumes \(Storage Groups\) list view](#) on page 335 and [z/OS unmap FBA volumes from the Volumes \(Storage Groups\) list view](#) on page 339.
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945 (Only available on storage systems running 5876).
- **Assign Symmetrix Priority**—[Assigning array priority to individual volumes](#) on page 189 (Only available on storage systems running 5876).
- **Pin**—[Pinning and unpinning volumes](#) on page 173 (Only available on storage systems running 5876).
- **Unpin**—[Pinning and unpinning volumes](#) on page 173 (Only available on storage systems running 5876).

Click  to view the Volume in Storage Group details view.

The following properties display:

- **Masking Info**—Number of masking views associated with the storage group.
- **Storage Groups**—Number of associated storage groups.

- **SRP**—Number of associated SRPs.
- **FBA Front End Paths** —Number of associated FBA Front End Paths.
- **RDF Info** —RDF Info.
- **Volume 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.
- **Encapsulated Device Flag** — Encapsulated device flag.
- **Encapsulated Device Array** — Encapsulated device array.
- **Encapsulated Device Name** — Encapsulated device name.
- **Status** — Volume status.
- **Reserved** — Whether the volume is reserved.
- **Capacity (GB)** —Volume capacity in GBs.
- **Capacity (MB)** —Volume capacity in MBs.
- **Capacity (CYL)** —Volume capacity in cylinders.
- **Compression Ratio** — Volume emulation.
- **Emulation** — Volume emulation.
- **AS400 Gatekeeper** — AS400 Gatekeeper indication.
- **Symmetrix ID** — Symmetrix system on which the volume resides.
- **Symmetrix Vol ID** — Symmetrix volume name/number.
- **HP Identifier Name** — User-defined volume name (1-128 alpha-numeric characters), applicable to HP-mapped devices. This value is mutually exclusive of the VMS ID.
- **VMS Identifier Name** — Numeric value (not to exceed 32766) with relevance to VMS systems. This value is mutually exclusive of the HP ID.
- **Nice Name** — Nice name generated by Symmetrix Engineuity.
- **WWN** — World Wide Name of the volume.
- **External Identity WWN** — External Identity 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 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).
- **Copy Pace - RDF** — Copy pace priority during RDF operations.
- **Copy Pace - Mirror Copy** — Copy pace priority during mirror operations.
- **Copy Pace - Clone** — Copy pace priority during clone operations.
- **Copy Pace - VLUN** — Copy pace priority during virtual LUN operations.
- **Dynamic Cache Partition Name** — Name of the cache partition.
- **Compressed Size (GB)** — Compressed size (GB).
- **Compressed Percentage** — Compressed percentage.
- **Compressed Size Per Pool (GB)** — Compressed Size Per Pool (GB).
- **XtremSW Cache Attached** — Indicates whether XtremSW cache is attached to the volume.
- **Base Address** — Base address.
- **AS400 Gatekeeper** — AS400 Gatekeeper indication.
- **Mobility ID Enabled** — Mobility ID enabled indication.
- **GCM** — GCM Flag set indication.

- **Optimized Read Miss** — Cacheless read miss status.
- **Persistent Allocation** — Persistent Allocation indication.
- **PowerPath Hosts** — Number of PowerPath hosts.
- **Mounted** — Mounted indication.
- **Process** — Process.
- **Last time used** — Last time used.

## Viewing Storage Group Compliance view

### Before you begin


The user requires a minimum of Monitor permissions to perform this task.

### Definitions:

- **Workload Skew** - Skew is represented by capacity and load pairs. There are two sources of skew for a storage group. One is using device stats. The other is using SG\_PER\_POOL chunks. There is an algorithm in WLP to merge these two lists to give us a usable skew profile. A skew profile is only useful if you have multiple chunks. If an SG has a single device, there is not enough data to calculate skew, the corresponding storage group per pool metrics can be used. Similarly, if an array has only one pool, the device stats are more meaningful for skew.
- **Workload Mixture** - The mixture is the distribution of various I/O types as percentages of the total IOPS. These are useful for determining, for example, whether a workload is heavy read or heavy write, whether I/Os are mostly random or mostly sequential.

To view the Storage Group (SG) Compliance view:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Groups View**.
3. Select a storage group and click  to view its details.
4. Select **VIEW ALL DETAILS**.
5. Select the **Compliance** tab.

Charts are displayed for the following:

- **Response Time chart** - this chart displays wait time weighted response time and (if applicable) the target service level response time band. The following section explains the data in the chart.
  - **Actual: running I/O to Storage Group** - Wait time weighted response time is calculated in buckets and displayed. If a bucket has no data, 0 is displayed.
  - **Actual: no I/O to Storage Group** - 0s are displayed.
  - **Planned: SLO Response Time Max and SLO Response Time Min** are displayed as a data band across the timeline. This is labeled "Planned". If the service level is Optimized, no plan is displayed, because there is no Response Time band for Optimized.
  - **Excluded Data:** If a recurring exclusion has been set via the Exclusion Windows dialog, the windows are represented by vertical gray plot bands.
  - **Last Processed:** A 2px dotted plot line marks the most recent SPA HOURLY timestamp processed by SPA for a given metric. It is not

represented in the legend, but if you hover, you can see the timestamp associated. In normal successful/processing, this acts as a "Where am I" indicator. If WLP stops processing for some reason, it's a subtle debugging helper.

- IOPS chart - This chart toggles between IO/sec and MB/sec, displaying IO rate weighted metric values, "planned" values, and (if set) Host IO Limits. The following section explains the data in the chart.
  - Actual: running I/O to Storage Group - IO Rate weighted total IOPS (or total MBPS) are calculated in buckets and displayed. If a bucket has no data, 0 is displayed.
  - Actual: no I/O to Storage Group - 0s are displayed.
  - Planned: Host I/O Limits for Standalone SG - Host IO Limit is displayed as a static value across the timeline. Host IO Limit is only shown on the chart it impacts. For example, if MBPS host IO limit is set, and the user has IOPS selected, they won't see anything unless they toggle to MBPS.
  - Planned: Host I/O Limits for Child SG, no limit for the parent SG - Host IO Limit is displayed as a static value across the timeline. Host IO Limit is only shown on the chart it impacts. For example, if MBPS host IO limit is set, and the user has IOPS selected, they won't see anything unless they toggle to MBPS.
  - Planned: No Host I/O Limits for Child SG and parent SG - If a cascaded SG has a host IO limit set at the parent, but no direct limit of its own, the host IO limit of any given child would be the parent limit minus whatever the siblings are using.
  - Planned: Host I/O Limits for Child SG and parent SG - If a cascaded SG has a host IO limit set at the parent, and a direct limit of its own, the host IO limit of any given child would be the more limiting of the parent limit minus whatever the siblings are using, or the child SGs own limit.
  - Excluded Data: If a recurring exclusion has been set via the Exclusion Windows dialog, the windows are represented by vertical gray plot bands.
  - Last Processed: A 2px dotted plot line marks the most recent SPA HOURLY timestamp processed by SPA for a given metric. It is not represented in the legend, but if you hover, you can see the timestamp associated. In normal successful/processing, this acts as a "Where am I" indicator. If WLP stops processing for some reason, it's a subtle debugging helper.
- Workload Skew chart - This chart compares actual workload skew - represented by cumulative capacity and load percentages (ordered by access density) - to planned skew. If there is no IO data, Actual is displayed as 50% skew - a straight line from (0,0) to (100,100). If there is one Device in SG AND Only One Thin Pool, then the merged device and sg per pool skew profile doesn't give us enough data points. Actual is displayed as 50% skew - a straight line from (0,0) to (100,100). If IO is running to the SG, the skew is a logarithmic curve (or stepped line graph in some cases).
- I/O Mixture chart - This chart compares actual workload mixture to planned workload mixture. The inner pie represents the actual IO distribution. The outer donut represents the planned mixture. If there is no I/O to the storage group, the mixture distribution will be equal percentages for each IO type (20% read hit, 20% sequential write, etc.) and the tooltip will show the corresponding IO sizes as 0kB.

Select the **Show Plan** slider to turn on or turn off the display of the plan. The plan is reference point used for comparison, and is a two week expiring performance reservation for subsequent provisioning suitability calculations.

The following controls are available:


- **Exclude Data** - [Managing Data Exclusion Windows](#) on page 158
- **Save As a Template** - [Creating storage templates](#) on page 267
- **Reset Workload Plan** - [Resetting Workload Plan](#) on page 177
- **Set Host I/O Limits** - [Setting host I/O limits](#) on page 132

## Viewing storage group performance details

### Before you begin

- The storage system is running HYPERMAX OS 5977 or higher.
- To perform this operation, a Monitor role is required.
- The storage system must be local and registered for performance.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Groups View**.
3. Select a storage group and click  to view its details.
4. Select **VIEW ALL DETAILS**.
5. Optional: Choose a day or time to retrieve performance related data.
6. Select the **Performance** tab.

Charts are displayed for the following:

- Read and Write response times
- Host MBs Read and Written per second
- Host reads and writes per second
- Read and Write response times
- FE Directors - Name, % Busy and queue depth utilization.
- FE Port - Name, % busy, and host I/Os per second.
- Related SGs - Name, response time, host I/Os per second, and host MBs per second.

## Select Storage Resource Pool

Use this dialog box to select a Storage Resource Pool for the operation.

---

### Note

To create the storage group outside of FAST control, set Storage Resource Pool to None; otherwise, leave this field set to the default.

---

## Select SSID

Use this dialog box to select an SSID for the operation.

## Task in Progress

Use this dialog box to monitor the progress of a configuration change operation

### Procedure

1. To view detailed information, click **Show Task Details**.

Once a task completes, a success or failure message displays.

## Select SRDF group

Use this dialog box to select a SRDF group.

## Editing storage group volume details

To edit storage group details for a storage system running Hypermax OS 5977 or higher:

### Procedure

1. Click the **Volume Config** tab.
2. 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 modifying storage groups with new volumes. Note that when modifying 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.
3. Optional: Click **Enable Compression** checkbox.
  4. Optional: Click the **Enable Mobility ID** checkbox to assign Mobility IDs to the volumes in the storage group. If you leave the checkbox unchecked, Compatibility IDs will be assigned to the volumes instead.
  5. Optional: Click **Allocate capacity for each volume** checkbox.
  6. Optional: Click **Persist preallocated capacity through reclaim or copy** checkbox.
  7. Click the **Volume Size** tab.

8. Enter a volume size, capacity and capacity unit.
9. Optional: Add one or more volume sizes by hovering over the area to the right of the volume capacity and selecting  .
10. Optional: Click  to remove a volume size.
11. Click **APPLY**.

The Storage Group page in the wizard displays *Mixed Capacities* for the row. Click *Mixed Capacities* to reopen this dialog.

## Editing storage group details

To edit storage group details for a storage system running Enginuity 5876:

### Procedure

1. 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.
- 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. (applicable for regular volumes only)

2. Click **OK**.

## Modify Custom Capacity dialog box

Use this dialog box to modify the capacity of a storage group with mixed capacities.

To modify the capacity of the storage group, type new values for the volumes and click **OK** to return to the **Modifying Storage Groups** dialog box.

# Understanding FAST

## Note

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

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

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

## Understanding service levels

A service level is the response time target for the storage group. The service level allows you set the storage 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 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 Levels




### Before you begin

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

Once a service level is renamed, all active management and reporting activities will be performed on the newly named service level. The original, pre-configured service level name will be maintained in the Service Level View for future reference. All other references to the original service level will display the new name.

### Procedure

1. Select the storage system.




2. Select **Storage > Service Levels** to open the **Service Level** view.
3. Hover over the service level name and click .
4. Type the new name over the existing name. Type the new name over the existing name and click  to complete the renaming process. To cancel the renaming, click .

## Reverting to original service level names

### Before you begin

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

### Procedure

1. Select the storage system.
2. Select **Storage > Service Levels** to open the **Service Level** view.
3. Hover over the service level name and click .
4. Type the original, pre-configured name. Type the new name over the existing name and click  to complete the renaming process. To cancel the renaming, click .

## Viewing service levels

### Before you begin

This feature requires HYPERMAX OS 5977 or higher.

A service level is the response time target for the storage group. The service level allows you set the storage 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 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.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Service Levels** to open the **Service Level** view.

For all-flash storage systems running HYPERMAX OS 5977, the only service level available is Diamond.

Available service levels are displayed in card format.

Each service level card shows the service level name (display name if it has been renamed), the expected average response time (in ms) and available headroom.



Clicking and selecting a card displays a table. The table gives more details for different workload types available for each service level. None workload type will be selected by default.


The details table is only visible for FBA service levels.

The columns in the table include workload type, target response time, headroom, I/O density, I/O size, write %, skew and usage count for each workload type. All columns are sortable and all columns (except workload type) can be hidden.

Hover near the write % to view a pop-out showing more details on the mixture and similarly hover near the skew % to view a pop-out showing more details about the skew.

3. Optional: To rename a service level, hover over the service level card and click



. Type the new name over the existing name and click  to complete the

renaming process. To cancel the renaming, click .

4. Optional: To provision storage using a service level, select the service level card and a corresponding workload type and click **Provision**.

This opens the **Provision Storage to Host** wizard, with the service level and the workload type will be populated by default. For CKD Provisioning wizard, only the service level will be selected by default. For more information on using the wizard, refer to [Using the Provision Storage wizard](#) on page 100.


## Changing service level

This functionality only applies to storage systems running HYPERMAX OS 5977 or higher and does not apply to all-flash storage systems.

For all-flash storage systems running HYPERMAX OS 5977 and higher, the only service level available is Diamond.

To change the service level:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Groups View**.
3. Select a storage group where performance data is available and click  to view its details.
4. Double click on the Compliance icon in the **Details** panel (alternatively, select **VIEW ALL DETAILS** in the **Details** panel and select the **Compliance** tab).
5. Click **VIEW DETAILS** (located at the top right of the **Response Time** panel).
6. Click **Change Service Level**.
7. Change the service level.
8. Click **OK**.

## 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.

### Procedure

1. Select the storage system.
2. Select **Storage > Storage Resource Pools** to open the **Storage Resource Pools** view.
3. Click **Modify**.
4. Modify any number of the following:
  - **Storage Resource Pool 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 **OK**.

## Viewing Storage Resource Pools

### Before you begin

- This feature requires HYPERMAX OS 5977 or higher.


### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Resource Pools** to open the **Storage Resource Pools** view.

The following properties display:

- **Name** — Name of the storage resource pool. .
- **Used Usable Capacity (%)**—Used usable capacity, expressed as a percentage.
- **Total Usable Capacity (TB)**—Total usable capacity.
- **Allocated Subscribed Capacity (%)**—Allocated subscribed capacity, expressed as a percentage.
- **Total Subscribed Capacity (TB)**—Total subscribed capacity.

The following controls are available:

- **Modify**—[Modifying Storage Resource Pool details](#) on page 154
- **Add EDisks**—[Adding external disks](#) on page 274
- Click  to view the following details:
  - **Name** — Name of the storage resource pool.
  - **Description** — Description.
  - **Default Emulation**—The default emulation for the pool (FBA or CKD).
  - **Overall Efficiency**—The current compression efficiency on this storage resource pool.
  - **Compression State** — Indicates whether compression is enabled or disabled for this storage resource pool.
  - **Effective Used Capacity (%)**—The effective used capacity, expressed as a percentage.
  - **Usable Capacity (TB)**—Usable capacity of all the disk groups in the Storage Resource Pool, excluding any external disk groups used for FTS encapsulation.
  - **Allocated Capacity (TB)**—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 (TB)** — 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.
  - **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.
  - **FBA Service Levels**—Number of FBA service levels..
  - **Disk Groups**—Number of Disk Groups.

The panel also provides links 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 the number next to **Disk Groups** will open a view listing the disk groups in the storage resource pool.

## Changing Storage Resource Pools for storage groups

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

In eNAS environments, you can also perform this operation from the File Storage Groups page (**System > System Dashboard > File Dashboard > File Storage Groups**).


Before you begin:

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

To change the Storage Resource Pool for storage groups:

### Procedure

1. Select the storage system.
2. Under **STORAGE**, select **Storage Groups**.
- 3.

Select the storage group, click , and select **Change SRP** to open the **Change SRP** dialog box.

4. Select the new SRP.
5. (Optional) Change the Service Level for the SG. Service levels 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 critical, 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.

For all-flash storage systems, the only service level available is Diamond and it is selected by default.

6. (Optional) Refine the service level by selecting the **Workload Type** to assign it. (This step is not applicable for storage systems running PowerMaxOS 5978.)
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 920 and [Previewing jobs](#) on page 920.

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

## Service level compliance

Each service level 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 service level. If both calculated values fall within (under) the service level 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 Compliance Reports

This procedure explains how to create Compliance reports. Compliance Reports allow you to view storage group performance against service levels over a period of time.

Before you begin:

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

To create Compliance Reports:

### Procedure

1. Select the storage system.
2. Select **SG COMPLIANCE**.
3. Within the **Compliance** panel, click **VIEW COMPLIANCE REPORT**.
4. Click **Schedule**.
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.
7. Optional: On the **Email** tab, select **Send report to** and type an email address.
8. Click **OK**.

## Viewing compliance reports

This procedure explains how to view storage group performance against service levels over a period of time.



Before you begin:

- This feature requires HYPERMAX OS 5977 or later.

- The user must be a StorageAdmin permissions or later.

To view service level compliance reports:

#### Procedure

1. Select the storage system.
2. Select **SG COMPLIANCE**.
3. Within the **Compliance** panel, click **VIEW COMPLIANCE REPORT**.
4. Customize the report by doing the following:
  - a. Select the time period. For the time period you select, the storage group's compliance is assessed in 30 minute intervals, and then its overall compliance state is displayed based on the method described in [Service level compliance](#) on page 157. For example, if you select **Last 24 hours**, the storage group's compliance state is assessed 48 times, and then its calculated compliance state is displayed in this report.
  - b.  Click  and select whether to view the compliance information as a chart or as numbers.

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 service level target.
- **% Marginal**—Percentage of time the storage group performed below the service level target.
- **% Critical**—Percentage of time the storage group performed well below the service level target.

The following controls are available:

- **Export**—Save the report as a PDF file.
- **Schedule**—[Creating Compliance Reports](#) on page 157.
- **Monitor**—[Performance Dashboards](#) on page 518

### Save Report Results dialog box

Use this dialog box to save service level compliance reports in PDF.

### Managing Data Exclusion Windows

This procedure explains how to manage Data Exclusion Windows for calculating headroom and suitability.

Peaks in storage system statistics can occur due to:

- anomalies or unusual events
- recurring maintenance during off-hours that fully loads the storage system

Due to the way this data is condensed and used, unexpected headroom and suitability results can occur.

There are two ways to improve the handling of these cases:

- **One-time exclusion period**—when the one-time exclusion period value is set, all statistics before this time are ignored. This helps resolve the first case above,

where a significant one time peak distorts the results due to reliance on two weeks of data points. This is set system-wide for all components.

- **Recurring exclusion period**—You can select one or more 4 hour windows to use in admissibility checks. This is set system-wide for all components. Recurring exclusion periods are repeating periods of selected weekday or time slot combinations where collected data is ignored for the purposes of compliance and admissibility considerations. The data is still collected and reported, but it is not used in those calculations.

Before you begin:

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

To manage Data Exclusion Windows:

### Procedure

1. Select the storage system.
2. Select **SG Compliance**.
3. In the **Actions** panel, select **EXCLUDE DATA**.

### Results

The **Compliance Settings** page allows you to view and set the one-time exclusion period and recurring exclusion periods for a selected storage system. It consists of two panels. The **One-time Exclusion Period** panel displays 84 component utilizations (two weeks worth of data) in a single chart that allows you to set the one-time exclusion period value from a given time slot, resulting in all time slots prior to the selected time slot being ignored for the purposes of calculating compliance and admissibility values. The **Recurring Exclusion Periods** panel displays the same data, but in a one-week format that allows you to select repeating recurring exclusion periods during which any collected data is ignored.

Each bar in the chart represents a utilization score calculated for that time slot. The score itself is the highest value of four component types, that is, the “worst performing” of the four components is the one that determines the overall value returned. The exact type and identifier of the selected component can be seen in the tool tip for a specific bar. The four component types that are represented in the bars are:

- Front End Port
- Back End Port
- RDF Port
- Thin Pool

The bars in both panels represent the same data using the same color coding scheme. The colors of the bars signify the following states:

Green represents a utilization value that meets the best practice limit.

Red represents a utilization value that exceeds the best practice limit.

Blue represents a utilization value this is being ignored before the one-time exclusion period.

Gray represents a utilization value that is being ignored as part of a window.

No Bar – If no data was collected or calculated during a time slot, there is no bar present.

The **One-time Exclusion Period** panel consists of a chart that is labeled with the component utilization value as the y-axis and the time slot as the x-axis. Each time slot

is a four hour window during which data was collected and a utilization score was calculated. There is also a horizontal line representing the best practice utilization of 100%. The x-axis is labeled with the dates of the time slots, that is, the dates of the midnight time slots are labeled with that date and other time slots are blank.

The top-right of this panel has a filter which allows you to include all components used in utilizations calculations or filter for only those used in headroom calculations. This can be helpful when headroom values are causing suitability problems in other areas, but those issues are masked by other component utilizations on this chart. The filters are: **All components, for suitability** (the default selection) and **Back-end components only, for headroom**. When you select a value, the page is reloaded with data from the server, filtered according to the selection made. Both charts are updated to reflect this data.

When the user selects a value the page will be reloaded with data from the server, filtered according to the selection made. Both charts will be updated to reflect this data.

You can select and set a time slot before which all collected data will be ignored. You select the time slot by clicking on the desired bar. The selected bar and all previous bars are changed to the one-time exclusion period coloring reflecting this selection. In addition, one-time exclusion period selection is also dynamically displayed in the Recurring Exclusion Periods chart as selections are made. If you try to set the selection to the last bar on the right an error is displayed and the action will not be allowed. In addition, selection is also dynamically displayed in the Windows chart as selections are made. You can deselect a selected bar by clicking it again. The chart then reverts to the value set when the page was loaded.

One-time exclusion period bars are only displayed in the **Recurring Exclusion Periods** chart under these conditions:

- Both buckets corresponding to the Recurring Exclusion Periods chart slot are before the one-time exclusion period.
- One of the buckets is before the one-time exclusion period and the other bucket has no data collected.

The panel has two buttons to set and clear any changes made:

- **Set One-time Exclusion**—writes the selected one-time exclusion period value to the database. This value will then be in effect and will be shown in all future views of this page. This button is enabled when a one-time exclusion period is selected. Clicking OK confirms the operation.
- **Clear One-time Exclusion**—clears any previously set one-time exclusion value. This button is only enabled if a one-time exclusion value is set when the page is first loaded. Clicking OK confirms the operation

The **Recurring Exclusion Periods** panel consists of seven charts, one for each day of the week. Each chart has a bar for each four hour time slot during which data is collected and a utilization score is calculated. Each bar represents two bars shown in the **One-time Exclusion Period Panel** chart. The bar shown in this chart is the highest value (“worst performing”) bar of the two **One-time Exclusion Period Panel** bars.

One-time exclusion period bars are only displayed in the **Recurring Exclusion Periods** chart under the following conditions:

- Both buckets corresponding to the Recurring Exclusion Periods chart slot are in the one-time exclusion period.
- One of the buckets is in the one-time exclusion period and the other bucket has no data collected.



In this panel, you click a time slot to select or deselect it. Clicking on a selected time slot will deselect it. As selections are made, both charts will be dynamically updated with the appropriate color coding.

The panel has two buttons to set and clear any changes made:

- **Set Recurring Exclusions**—writes the selected recurring exclusions period value(s) to the database. These values will then be in effect and will be shown in all future views of this page. This button is enabled when a recurring exclusion period is selected. Clicking OK confirms the operation.
- **Clear Recurring Exclusions**—clears any previously set recurring exclusion period values. This button is only enabled if a recurring exclusion period value is set when the page is first loaded. Clicking OK confirms the operation.

At the bottom of the page is a panel that contains the legend indicating the meanings associated with the different bar colors. On the right hand side in this panel is text detailing the last time a one-time exclusion period or Window was changed. If you hover over this text, the name of the user (fully qualified user name) that performed the last update operation is displayed. If the database has never had a one-time exclusion period or Window set, the field and tool tip text displays “Not yet modified”.

#### Alerts

There is a system alert generated each time a user changes a one-time exclusion period value or a recurring exclusion period value.

## Symmetrix tiers

### Creating tiers

#### Before you begin

- This feature requires Enginuity 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 tier:

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Tiers** to open the Tiers list view.
3. Click **Create** to open the **Create Tier** dialog box.

When this dialog box first opens, the chart displays the configured and unconfigured space on the selected storage system. Once you select a disk group or thin pool, this chart displays 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 storage 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. Optional: Select **Include all future disk groups on matching technology** for this tier.
- Tiers created in this manner are considered dynamic tiers. Tiers created without this option are considered static tiers.
12. Click **OK**.

## Modifying tiers

### Before you begin

- This feature requires Enginuity 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 (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 tier:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Tiers** to open the **Tiers** list view.

3. Select the tier and click **Modify**.
4. Add or remove disk groups/thin pools by selecting/clearing the corresponding check box.
5. Click **OK**.

## Renaming tiers



### Before you begin

- This feature requires Enginuity 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 storage system (across both DP and VP tier types), ignoring differences in case.

To rename a tier:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Tiers** to open the **Tiers** list view.
3.
 


 Select the tier, click , and click **Rename**.
4. Type a new name for the tier.
5. Click **OK**.

## Deleting tiers


### Before you begin

- This feature requires Enginuity 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 tier:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Tiers** to open the **Tiers** list view.
3.
 

Select the tier and click .
4. Click **OK**.

## Viewing Symmetrix tiers

### Before you begin

This feature requires Enginuity 5876.

### Procedure



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.
- **Protection**—RAID protection level assigned to the volumes 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.

The following controls are available:

-  — [Viewing Symmetrix tier details](#) on page 164
- **Create** — [Creating tiers](#) on page 161
- **Modify** — [Modifying tiers](#) on page 162
-  — [Deleting tiers](#) on page 163
- **Rename** — [Renaming tiers](#) on page 163


## Viewing Symmetrix tier details

### Before you begin

This feature requires Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Tiers** to open the **Tiers** list view.
- 3.

Select the tier and click  to open its **Tier Demand Report** panel or its **Details** panel.

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

The following properties display in the **Details** panel:


- **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.  
**Disk Location**—Internal or external.
- **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.
- **Number of Thin Pools** — Number of Thin Pools. Clicking the number next to **Number of Thin Pools** opens a view listing the associated thin pools.

## Viewing thin pools in a storage tier

### Procedure


1. Select the storage system.
2. Select **Storage > Tiers** to open the **Tiers** list view.
3. Select the tier and click  to open its **Details** view.
4. Click the number next to **Number of Thin Pools** to open the tier's **Thin Pool** view.

This view allows you to view and manage a tier's thin pool.

The following properties display:

- **Name**—Pool name.
- **Technology**—Disk technology type.
- **Configuration**—Protection configuration.
- **Emulation**—Pool emulation type based on the first volume added to the pool .
- **Allocated Capacity** — Percent capacity allocated to the pool .
- **Enabled Capacity (GB)**—Pool capacity in Gigabytes.

The following controls are available:

-  —[Viewing thin pool details](#) on page 246
- **Create**—[Creating thin pools](#) on page 240
- **Modify**—
- **Expand**—[Expanding thin pools](#) on page 241
- **Delete**—[Deleting thin pools](#) on page 243
- **Start Write Balancing**—[Starting and stopping thin pool write balancing](#) on page 242
- **Stop Write Balancing**—[Starting and stopping thin pool write balancing](#) on page 242
- **Bind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257

## FAST policies

### Creating FAST policies

#### Before you begin

- This feature requires Enginuity 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; however, you can create such policies in Solutions Enabler. The *Solutions Enabler Array Controls and Management CLI User 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.

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.

#### Procedure

1. Select the storage system.
2. Select **STORAGE > FAST Policies**.
3. Click **Create**.
4. 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.
5. Select the host type.
6. Select the volume **Emulation**.
7. 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.
8. Repeat the previous step for any additional tiers you want to add.
9. Click **OK**.

## Modifying FAST policies

#### Before you begin

- This feature requires Enginuity 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.

**Procedure**

1. Select the storage system.
2. Select **STORAGE > FAST Policies**.
3. Select a policy and click **Modify**.
4. Optional: Modify the **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.
5. Optional: Change the host type.
6. Optional: Change the volume **Emulation**.
7. Optional: Select a tier to modify for 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.
8. Repeat the previous step for any additional tiers you want to modify.
9. Click **OK**.

**Deleting FAST policies****Before you begin**

- This feature requires Enginuity 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:

**Procedure**

1. Select the storage system.
2. Select **STORAGE > FAST Policies**.
3. Select the policy and click **Delete**.
4. Click **OK**.

**Associating FAST policies with storage groups****Before you begin**

Storage groups and FAST policies can only be associated under the following conditions:

- The storage system is 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 storage system has fewer than the maximum number of allowed associations (8,192).

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.

To associate a FAST policy with a storage group:

#### Procedure

1. Select the storage system.
2. Select **STORAGE > FAST Policies**.
3. Select the policy and click **Associate Storage Groups**.
4. Select one or more storage groups to be associated with the FAST policy.
5. To have FAST factor the R1 volume statistics into move decisions made for the R2 volume, select **Enable FAST VP 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. Both R1 volumes and R2 volumes need to be running Enginuity version 5876 or higher for the FAST VP system to coordinate the moves.

6. Click **OK**

## Associating storage groups with FAST policies

### Before you begin

Storage groups and FAST policies can only be associated under the following conditions:



- The storage system is 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 storage system has fewer than the maximum number of allowed associations (8,192).

To associate a storage group with a FAST policy:

#### Procedure



1. Select the storage system.
2. Select **STORAGE > Storage Groups**.
3.
 


 Select the storage group, click , and select **FAST > Associate**.
4. Select a policy and click **OK**.

### Disassociating FAST policies and storage groups

#### Procedure

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 **FAST > Disassociate**.
4. Click **OK**.

### Reassociating FAST policies and storage groups

#### Before you begin


- This feature requires Enginuity 5876.
- The storage group name must be valid.
- The storage group and policy must already exist on the storage 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.


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.

#### Procedure

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 **FAST > Reassociate**.
4. Select a policy and click **OK**.

## Viewing FAST policies

#### Before you begin

This feature requires Engenuity 5876.

#### Procedure


1. Select the storage system.
2. Select **STORAGE > FAST Policies**.

Use the **FAST Policies** list view to view and manage FAST policies on a storage system.

The following properties display:

- **Name**—Name of the policy.
- **Type**—Type of the policy.
- **Tier 1**—Storage tier associated with the policy.
- **Tier 2**—Storage tier associated with the policy.
- **Tier 3**—Storage tier associated with the policy.
- **Tier 4**—Storage tier associated with the policy.
- Up to 4 tiers is supported only for FAST VP policies. FAST policies support up to 3 tiers.
- **Storage Groups**—Storage groups associated with the policy.


The following controls are available:

-  —[Viewing FAST policy details](#) on page 172
- **Create**—[Creating FAST policies](#) on page 166
- **Modify**—
- **Delete**—[Deleting FAST policies](#) on page 168
- **Associate Storage Group**—[Associating storage groups with FAST policies](#) on page 169

## Viewing storage group for FAST policies

#### Procedure

1. Select the storage system.
2. Select **STORAGE > FAST Policies**.

3. Select the policy, click , and select the **Details** panel.
4. Click the number next to **Storage Groups**.

The following properties display:

- **Name**—Name of the storage group.
- **FAST Policy**—Policy associated with the storage group.
- **Capacity**—Total capacity of the storage group in GB.
- **Volumes**—Number of volumes contained in the storage group.
- **Masking Views**—Number of masking views associated with the storage group.

Refer to [Viewing storage groups](#) on page 135 for information on properties and controls for the storage group.


## Viewing FAST policy details

### Before you begin

This feature requires Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > FAST Policies**.
- 3.

Select the policy and click  to open its **Tier Demand Report** panel or its **Details** panel.

The **Tier Demand Report** 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.

In the **Details** panel, the following properties display:

- **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.
- **Storage Groups** — Number of Storage Groups. Clicking the number next to **Storage Groups** will open a view listing the associated storage groups.

- **Number of Tiers** — Number of Tiers. Clicking the number next to **Number of Tiers** will open a view listing the tiers in the policy.

## Viewing FAST storage groups

This feature is only supported on storage systems running Enginuity OS 5876.

Up to 4 tiers is supported only for FAST VP policies. FAST policies support up to 3 tiers.

### Procedure

1. Select the storage system.
2. Select **SG COMPLIANCE** on the dashboard.
3. Select **VIEW FAST STORAGE GROUPS**.

The following properties display:

- **Name**—Name of the storage group.
- **FAST Policy**—FAST Policy associated with the storage group.
- **Compliant**—Icon indicating compliance.
- **Tier 1 %**—Storage tier percentage associated with the policy.
- **Tier 2 %**—Storage tier percentage associated with the policy.
- **Tier 3 %**—Storage tier percentage associated with the policy.
- **Tier 4 %**—Storage tier percentage associated with the policy.
- **Out of Policy %**—Out of Policy percentage.

## Pinning and unpinning volumes

### Before you begin

This feature requires Enginuity 5876.


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

### Note

The capacity of pinned volumes is counted for compliance purposes.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Select the volume type by selecting a tab.
- 4.

Select one or more volumes, click  , and select one of the following:

- **FAST > Pin**—To pin the volumes.
- **FAST > Unpin** —To unpin the volumes.

5. Click **OK**.

## 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.

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

#### Before you begin

- This feature requires Enginuity OS 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 storage system is 128.

#### Procedure

1. To create time windows:
  1. Select the storage system.
  2. Navigate to **DASHBOARD > SG COMPLIANCE**.
  3. From within the **FAST Status Report** panel, select FAST VP or FAST DP (If the storage system is licensed for both FAST DP and FAST VP) for which the time window will apply.
  4. Click  next to the type of time window you want to create or modify.  
Depending on your selection, either the **Performance Time Window** or the **Move Time Window** dialog opens.
  5. If you are creating or modifying an open time window, select the day(s) or week in which to define the time window and click **ADD**.
  6. Select one of the following options:
    - **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.
7. Click **OK**.
  8. If you are creating or modifying a closed time window, select the **Start Time** checkbox and click **ADD**.
  9. Select the start date and time and the end date and time and click **OK**.
  10. 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.
  11. Click **SAVE**.

## Deleting time windows

### Before you begin

Time windows are used by FAST.

### Procedure

1. To delete time windows:
  1. Select the storage system.
  2. Navigate to **DASHBOARD > SG COMPLIANCE**.
  3. From within the **FAST Status Report** panel, select FAST VP or FAST DP (If the storage system is licensed for both FAST DP and FAST VP) for which the time window will apply.
  4. Click  next to the type of time window you want to create or modify.  
Depending on your selection, either the **Performance Time Window** or the **Move Time Window** dialog opens.
  5. If you are deleting an open time window, select the day(s) and click **REMOVE**.
  6. If you are deleting a closed time window, select the **Start Time** checkbox and click **REMOVE**.

## FAST Movement Time Window dialog box

Use this dialog box to manage movement time windows, including the following tasks:

- [Creating and modifying time windows](#) on page 174
- [Deleting time windows](#) on page 175

## FAST Performance Time Window dialog box

Use this dialog box to manage performance time windows, including the following tasks:

- [Creating and modifying time windows](#) on page 174
- [Deleting time windows](#) on page 175

## Manage Closed Movement Time Windows dialog box

Use this dialog box to manage closed movement time windows, including the following tasks:

- [Creating and modifying time windows](#) on page 174
- [Deleting time windows](#) on page 175

## Manage Closed Performance Time Windows dialog box

Use this dialog box to manage closed movement time windows, including the following tasks:

- [Creating and modifying time windows](#) on page 174
- [Deleting time windows](#) on page 175

## Manage Open Movement Time Windows dialog box

Use this dialog box to manage open movement time windows, including the following tasks:

- [Creating and modifying time windows](#) on page 174
- [Deleting time windows](#) on page 175

## Manage Open Performance Time Windows dialog box

Use this dialog box to manage open performance time windows, including the following tasks:

- [Creating and modifying time windows](#) on page 174
- [Deleting time windows](#) on page 175

## 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.

Workload Planner is supported on storage systems running Enginuity 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 592 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.
  - For storage systems running Enginuity 5876, they must be associated with a FAST policy.
- Contain only FBA volumes.

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

## Delete a reference workload

This dialog allows you to delete a reference workload. Click **OK** to confirm.

## Resetting Workload Plan

### Before you begin

To perform this operation, a StorageAdmin role is required.

Resetting the workload plan requires one week of data.

This procedure explains how to set the performance baseline expectations of a storage group to the characteristics currently measured for the previous two weeks.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**.
3. Select the storage group and click the Compliance icon to open its details view.
4. If not already displaying, click the **Compliance** tab.
5. Click **Reset Workload Plan**.
6. Review the **Current Scores** and the projected **New Baseline**.
7. If satisfied, click **OK**.

### Results

Once complete, the **Workload Planning** tab updates with the newly calculated performance metrics.

# Managing volumes

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 5876, refer to [Managing volumes](#) on page 178.

To view volumes associated with a host initiator, refer to [Viewing volumes associated with host initiator](#) on page 315.

---

To use the **Volumes** view:

**Procedure**

1. Select the storage system.
2. Select **Storage > Volumes** to open the **Volumes** list view.

For field and control descriptions, refer to the following volume-specific help pages:

- TDEV—[Viewing thin volumes](#) on page 223
- DATA—[Viewing DATA volumes](#) on page 239
- CKD—[Viewing CKD volumes](#) on page 92

## Managing volumes

For storage systems running Enginuity OS version 5876, the Volumes view provides you with a single place from which to view and manage all the volume types on the system.

---

**Note**

For instructions on managing volumes on storage systems running HYPERMAX OS 5977, refer to [Managing volumes](#) on page 177.

---

To use the Volumes view:

**Procedure**

1. Select the storage system.
2. Select **Storage > Volumes** to open the Volumes view.

The Regular Volumes list view is displayed by default (see [Viewing regular volumes](#) on page 217).

Click the **Virtual** tab to see the Virtual Volumes list view (see [Viewing virtual volumes](#) on page 224).

Click the **Meta** tab to see the Virtual Volumes list view (see [Viewing meta volumes](#) on page 208).

Click the **Private** tab to see the Virtual Volumes list view (see [Viewing private volumes](#) on page 215).

## Creating volumes

This procedure explains how to create volumes.

**Procedure**

1. Select the storage system.
2. Select **STORAGE > Volumes** and click **Create** 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](#) on page 185
    - CKD—[Creating CKD volumes](#) on page 330

- **Virtual Gatekeeper**—[Creating virtual gatekeeper volumes](#) on page 186

---

#### Note

The maximum volume size supported on a storage system running HYPERMAX OS 5977 or higher is 64 TB.

---

- Enginuity 5876:
  - **DATA**—[Creating DATA volumes](#) on page 179
  - **Diskless**—[Creating diskless volumes](#) on page 180
  - **DRV**—[Creating DRV volumes](#) on page 181
  - **Gatekeeper**—[Creating gatekeeper volumes](#) on page 181
  - **Regular**—[Creating regular volumes](#) on page 182
  - **SAVE**—[Creating SAVE volumes](#) on page 183
  - **TDEV**—[Creating thin volumes](#) on page 184
  - **VDEV**—[Creating VDEV volumes](#) on page 187

In addition, you can also create volumes using a storage template.

## Creating DATA volumes

This procedure explains how to create DATA volumes on storage systems running Enginuity version 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click ⓘ to open its **Details** view.
4. Click the number next to **Number of Data Volumes**.
5. Click **Create Volumes**.
6. Select **DATA** as the **Configuration**.
7. Select the **Disk Technology**.
 

**External** disk technology is an option if the storage system has FTS (Federated Tiered Storage) enabled and available external storage.
8. Select the **Emulation** type.
9. Select the RAID **Protection** level.
10. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
11. To add the new volumes to a specific thin pool, select one from **Add to Pool**. Pools listed are filtered on technology, emulation, and protection type.
12. Click **Advanced Options** to continue setting the advanced options, as described next.

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 the 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**.
  - d. Click **APPLY**.
13. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Creating private volumes

The following private volumes can be created:

[Creating DATA volumes](#) on page 179

[Creating diskless volumes](#) on page 180

[Creating DRV volumes](#) on page 181

[Creating gatekeeper volumes](#) on page 181

[Creating SAVE volumes](#) on page 183

## Creating diskless volumes

This procedure explains how to create diskless volumes on storage systems running Enginuity version 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** tab.
3. Filter on Type and select **DLDEV**.
4. Click **Create**.
5. Select the **Configuration** type.
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 the volume capacity.
8. To add the new volumes, select one from **Add to Pool**.
9. Click **Advanced Options** to continue setting the advanced options, as described next.

Setting advanced options

- a. Modify the **Volume Identifier**.
- b. 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.

- c. 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.
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 920 and [Previewing jobs](#) on page 920.
  - 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 storage systems running Enginuity version 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Private** tab.
3. Filter on Type and select **DRV**.
4. Select **DRV** as the **Configuration**.
5. Click **Create**.
6. Select the **Configuration** type.
7. Select the **Emulation** type.
8. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
9. To add the new volumes, select one from **Add to Pool**.
10. Click **Advanced Options** to continue setting the advanced options, as described next.

Setting Advanced options:

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.

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 920 and [Previewing jobs](#) on page 920.
  - 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 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Regular** panel.
3. Select the volume and click **Create**.
4. Select **Gatekeeper** as the **Configuration**.
5. Select the **Emulation** type.
6. Type the **Number of Volumes** to create.
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 920 and [Previewing jobs](#) on page 920.
  - 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 Engenuity version 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Regular** panel.
3. Click **Create**.
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 **Advanced Options** to continue setting the advanced options, as described next.

Setting Advanced options:

- a. 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).
- b. To create the volumes from a specific **Disk Group**, select one (disk group number and name).
- c. 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 196.

- d. 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.
  - e. 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.
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 920 and [Previewing jobs](#) on page 920.
    - 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 storage systems running Enginuity version 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Private** tab.
3. Filter on Type and select **SAVE**.
4. Select **SAVE** as the **Configuration**.
5. Select the **Disk Technology**.
6. **External** disk technology is an option if the Symmetrix system has FTS (Federated Tiered Storage) enabled and available external storage.
7. Select the **Emulation** type.
8. Select the RAID **Protection** level.
9. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.
10. 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.

## 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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Click **Advanced Options** 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:

- 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**. If Auto meta is enabled on the system then it displays as enabled with a green check mark.
- 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Creating thin volumes

This procedure explains how to create thin volumes on storage systems running Enginuity version 5876. For instructions on creating thin volumes on storage systems running HYPERMAX OS 5977 or higher, refer to [Creating thin volumes](#) on page 185.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**, click on the **Virtual** tab and select **Create**.
3. Select **Configuration (TDEV or BCV + TDEV or Virtual Gatekeeper)**.
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. 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).
7. Click **Advanced Options** to continue setting the advanced options

Setting Advanced options:

- a. 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 196.

- b. To **Allocate Full Volume Capacity**, select the option.
- c. 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.
- d. 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.
- e. 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.

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 920 and [Previewing jobs](#) on page 920.
  - Expand Add to Job List, and click **Run Now** to perform the operation now.
  - Click **Advanced Options** to continue setting the advanced options, as described next.

## Creating thin volumes

This procedure explains how to create thin volumes on storage systems running HYPERMAX OS 5977. For instructions on creating thin volumes on storage systems running Enginuity 5876, refer to [Creating thin volumes](#) on page 184.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click **Create** 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 storage group, and then click **OK**.
7. Click **Advanced Options** to set the advanced options:
  - Optional: Click the **Enable Mobility ID** checkbox to assign Mobility IDs to the volume. If you leave the checkbox unchecked, a Compatibility ID will be assigned to the volume instead.
  - 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.
  - 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 196.

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 920 and [Previewing jobs](#) on page 920.
  - 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.

This procedure explains how to create virtual gatekeeper volumes.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click **Create** to open the **Create Volume** dialog box.
3. Select **Virtual Gatekeeper** as the **Configuration**.

4. Optional: Select the **Emulation** type.
5. Type the **Number of Volumes**.
6. Optional: To add the volumes to a storage group, select the storage group 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Creating VDEV volumes

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** tab.
3. Filter on Type and select **VDEV**.
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. 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).
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
  - Click **Advanced Options** to continue setting the advanced options, as described next.  
Setting Advanced options:
    - a. View **Enable SCSI3 Persistent Reservation** status — For Engenuity 5876 and higher this feature is pre-set by SYMAPI and cannot be changed. It is displayed as enabled for Engenuity 5876 and higher, except for CDK and AS/400 emulations.
    - b. 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.
    - 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 920 and [Previewing jobs](#) on page 920.

- Expand **Add to Job List** and click **Run Now** to create the volumes now.


## Select Storage Group

Use this dialog box to select a storage group for the operation.

## Deleting volumes

This procedure explains how to delete volumes.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume that you wish to delete.
4. Select the volume and click .

## Duplicating 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.

The following explains how to duplicate volumes.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume that you wish to duplicate.
4. Select the volume, click , and click **Configuration > Duplicate Volume**.
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 **Advanced Options** to continue setting the advanced options.

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 196.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List** and click **Run Now** to create the volumes now.

## Assigning array priority to individual volumes



### Before you begin

This feature requires Enginuity 5876.

This procedure explains how to prioritize the service time of the host I/O to an individual volume. To prioritize the service time of the host I/O to groups of volumes (device groups or storage groups), refer to [Assigning array priority to groups of volumes](#) on page 189.

To assign host priority to individual volumes:

### Procedure

1. Select the storage system.
2. Select **Storage > Volumes**.
3. Click on the appropriate volume panel.
4.  Select the volume, click , and select **Assign Symmetrix Priority**.
5. Select an array priority from 1 (the fastest) to 16 (the slowest) and click **OK**.

## Assigning array priority to groups of volumes



### Before you begin

This feature requires Enginuity 5876.

This procedure explains how to prioritize the service time of the host I/O to groups of volumes (device groups or storage groups).

### Procedure

1. Select the storage system.
2. Do one of the following:
  - To assign priority to storage groups, select **STORAGE > Storage Groups** to open the Storage Groups list view .

- - Select the storage group, click , and select **Assign Symmetrix Priority** to open the **Assign Symmetrix Priority** dialog box.
  - ■ To assign priority to device groups, select **DATA PROTECTION > Device Groups** to open the Device Groups list view .
  - - Select the device group, click , and select **Assign Symmetrix Priority** to open the **Assign Symmetrix Priority** dialog box.
3. Select an array priority from 1 (the fastest) to 16 (the slowest) and click **OK**.
  4. Click **OK**.

## Changing volume configuration

### Before you begin

- On storage systems running Enginuity 5876 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 RAID set 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.

This procedure explains how to change a volume's configuration.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume.
- 4.

Select the volume, click , and click **Configuration > Change Volume Configuration**.

5. 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.
6. 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 storage systems with ESCON or FICON directors (or mixed systems). This field is optional on storage systems running Enginuity 5876 or higher when reducing the number of mirrors.

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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List** and click **Run Now** to run the job now.


## Expanding existing volumes

### Before you begin

- Requires HYPERMAX OS 5977 or later (HYPERMAX OS 5977.1125.1125 or later for CKD volumes).
- You must be logged in as an Administrator.
- You can expand a volume up to 64 TB (for FBA volumes) or 1,182,006 cylinders (for CKD volumes).
- When expanding a CKD volume above 565,250 cylinders, the new size must be a multiple of 1113 cylinders. If you specify an amount that is not a multiple, the system rounds it up.
- Consider consulting with your operating system vendor or cluster vendor for support of online LUN expansion
- You cannot expand a FBA volume when any of the following operations are in progress:
  - Free all
  - Reclaim
  - Deallocation
- Restrictions apply when a volume:
  - is a gatekeeper
  - is an ACLX
  - is Celerra FBA
  - is AS400
  - is VP encapsulated
  - is part of a SnapVX session defined
  - is being replicated
  - is part of an SRDF pair
  - is part of an ORS session
  - a TDAT
- For CKD volumes, you cannot expand a volume that is:
  - A CKD 3380 volume
  - Marked as Soft Fenced

### Procedure

1. Select the storage system.
- 2.

Select **STORAGE > Volumes**, click , and click **Expand Volume** to open the **Expand Volume** dialog box.

The **Expand Volume** dialog box appears.

3. 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.


SRDF storage group volume capacity can be expanded using the controls. In the case of SRDF Storage Group volumes, you need to specify a SRDF group number so that the dialog allowing you to remote volumes can also be displayed (see [Expanding remote volumes](#) on page 511).

4. To reserve the volume, select **Reserve Volumes**.
5. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Mapping volumes

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume.
- 4.

Select the volume, click , and click **Configuration > Map**.

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 the 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:
  - **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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Unmapping volumes

This procedure explains how to unmap volumes.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume.
4.
 

Select the volume, click , and click **Configuration > Unmap**.
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 [Setting preferences](#) on page 49 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. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting optimized read miss

### 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**.

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.

**Procedure**

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume.
- 4.

Select the volume, click , and select **Set Optimized Read Miss**.

5. 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 configuration.
  - **On**—Enables optimized read miss mode for both XtremCache and non-XtremCache EFD-only configurations.
6. Click **OK**.

**Setting volume status****Before you begin**

You cannot set the status of an unbound thin volume.

To set volume status for individual volumes:

**Procedure**

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume.
- 4.

Select the volume, click , and click **Set Volumes > Status**.

5. Set the 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.
6. Optional: For HYPERMAX OS 5977 or higher, select **SRDF/Metro**.
7. Optional: To force the operation when the operation would normally be rejected, select **SymForce**, if available.

8. If the selected volumes are mapped, you can select to change the status for a particular **Director** or all directors.
9. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting volume attributes

### Before you begin

You cannot set attributes for DATA volumes.


Setting attributes for CKD volumes is not supported. If attempting to set attributes for multiple volumes of type FBA and CKD, a warning is displayed stating that the action will be applied only to FBA volumes.

Setting the volume attribute for a volume restricts how it can be accessed.

To set volume attributes:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume.
- 4.

Select a volume, click , and click **Set Volumes > Attribute**.

5. 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. This feature only applies/appears for storage systems running Enginuity 5876.
  - **Dynamic RDF Capability** — Sets the volume to perform dynamic RDF operations. This feature only applies/appears for storage systems running Enginuity 5876. Possible operations are:
    - **No Change** — Keeps the RDF capability the same.
    - **Dynamic RDF Capability** — Sets the volume to perform dynamic RDF operations. This feature only applies/appears for Symmetrix systems running Enginuity 5876. Possible operations are:
    - **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** — This can be set to enabled or disabled. Maintains any reservations (flags) whether the system goes online or offline. This field will appear dimmed for diskless 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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting volume identifiers

This operation can be invoked from multiple locations in the Unisphere user interface. Depending on where the operation is invoked, some of the steps below may not apply.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume.
- 4.

Select the volume, click , and click **Set Volumes > Identifier**.

5. Set the volume identifiers:
  - 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.
  - Type the **Volume HP Identifier Name**. HP identifier names must be a user-defined volume name (not to exceed 64 alpha-numeric characters and underscores ( \_ ) ) applicable to HP-mapped volumes. This value is mutually exclusive of the VMS ID. This attribute will appear grayed out for diskless volumes.
  - 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.
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 920 and [Previewing jobs](#) on page 920.
  - 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 copy pace (QoS) for device groups

### Procedure

1. Select the storage system.
2. Select **Data Protection > Device Groups**.
3.  Select the device group, click , and select **Replication QoS**.
4. Select **Operation Type** from the following valid values:
  - **SRDF** — Sets the copy pace priority during SRDF 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 only available on arrays running Enginuity 5876 or higher.
5. Select the **Copy Pace** from the following valid value:
  - **0 -16** — Sets the copy pace, with 0 (the default) as the fastest and 16 as the slowest.
  - **STOP** — Stops the copy. Not supported when the **Operation Type** is BCV, or the array is running an Enginuity version lower than 5876.
  - **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 array is running an Enginuity version lower than 5876.
6. If performing this operation on a group: Select the type of volumes on which to perform the operation.
7. Click **OK**.

## QOS for replication

The QoS (Quality of Service) feature adjusts 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, the overall performance of other storage volumes increases.




The following tasks are supported:

- [Setting copy pace \(QoS\) for storage groups](#) on page 197
- [Setting copy pace \(QoS\) for device groups](#) on page 197
- [Setting copy pace \(QoS\) for volumes](#) on page 198

## Setting copy pace (QoS) for storage groups


### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Groups** list view.
3. Perform one of the following actions:

- For all volumes in the storage group: Select the storage group, click , and select **Replication QoS** to open the **Set Replication Priority QoS** dialog box.
- For some volumes in the storage group:
  - a. Select a storage group.
  - b. Click , and select the number next to **Volumes** to open the **Volumes** list view.
  - c. Select the volumes, click , and select **Set Volumes > Replication QoS**.
- 4. Optional: Select the **Show Selected Group Copy Pace** checkbox.
- 5. Select **Operation Type** from the following valid values:
  - **SRDF** — 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 only available on arrays running Enginuity 5876 or higher.
- 6. Select the **Copy Pace** from the following valid values:
  - **0 -16** — Sets the copy pace, with 0 (the default) as the fastest and 16 as the slowest.
  - **STOP** — Stops the copy. Not supported when the **Operation Type** is BCV, or the array is running an Enginuity version lower than 5876.
  - **URGENT** — Sets the copy pace to urgent, which may be faster than the default (0). Not supported when the **Operation Type** is BCV.
- 7. Click **OK**.

## Setting copy pace (QoS) for volumes

### Procedure

1. Select the storage system.
2. Select **Storage > Volumes**.
3. Navigate to the volume.
4. Select one or more volumes, click , and select **Replication QoS**.
5. Select **Operation Type** from the following valid values:
  - **SRDF** — 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 only available on arrays running Enginuity 5876 or higher.
6. Select the **Copy Pace** from the following valid values:

- **0 -16** — Sets the copy pace, with 0 (the default) as the fastest and 16 as the slowest.
- **STOP** — Stops the copy. Not supported when the **Operation Type** is BCV, or the array is running an Enginuity version lower than 5876.
- **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 array is running an Enginuity version earlier than 5876.

7. Click **OK**.

## Managing Meta Volumes

### Creating meta volumes

#### Before you begin

- Meta volumes are supported on storage systems running Enginuity 5876.
- 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 and Management CLI User Guide*.
- When creating meta volumes, 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.

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Meta** tab.
3. Click **Create**.
4. Select the **Emulation** type.
5. If creating FBA volumes, select whether to create them from **Create Volumes** or **Use Existing Volumes** volumes.
6. If creating FBA or AS/400 volumes, select the **Meta Volume Configuration (Concatenated or striped)**.
7. Select a method for forming the meta volumes.
8. Click **Next**.
9. 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 [Setting preferences](#) on page 49 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 **Advanced Options**, 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**.
- 10. Verify your selections in the **Summary** page. To change any of your selections, click **Back**.  
 Note that some changes may require you to make additional changes to your configuration.
- 11. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Adding meta members

### Before you begin

- Meta volumes are supported on storage systems running Enginuity 5876.
- To expand a bound striped thin meta volume on a storage system running Enginuity 5876 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 244.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Meta** tab.
3. Select the meta volume and click **Add Member**.
4. 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.
5. Select one or more volumes to add to the meta volume.
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 920 and [Previewing jobs](#) on page 920.


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

## Removing meta members

### Before you begin

- Meta volumes are supported on storage systems running Enginuity 5876.
- You can only remove members from concatenated meta volumes.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Meta** tab.
3. Select the meta volume and click  to open its **Details** view.
4. Click the number next to **META Members** to open the **Meta Members** list view.
5. Select one or more members and click **Remove Meta Member** to open the **Remove Meta Volume Member** dialog box.
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 920 and [Previewing jobs](#) on page 920.
  - 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 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Meta** tab.
3. Select the meta volume and click **Dissolve**.
4. Optional: If required, select **Delete Meta Members after dissolve**.  
Note that selecting **Delete meta members after dissolve** requires the operation to be run immediately (it cannot be scheduled).
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Converting meta volumes

### Before you begin

Meta volumes are supported on storage systems running Enginuity 5876.

This procedure explains how to change the configuration of a meta volume.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Meta** tab.
3. Select the meta volume and click **Convert**.
4. 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.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing CKD volumes

See below for procedure to view CKD volumes from the **Hosts > Mainframe** dashboard. To see the CKD volumes in a CU image, see [Viewing CKD volumes in CU image](#) on page 93.

### Procedure


1. Select the storage system.
2. Select **HOSTS > Mainframe** and click on **CKD Volumes** in the **Summary** panel.


The **CKD Volumes** list view is displayed. Use the this list view to view and manage the volumes.

The following properties display: however, not all properties may be available for every volume type:

- **Name**—Assigned volume name.
- **Type**—Type of volume.
- **Allocated %**—% of the volume that is allocated.
- **Capacity (GB)**—Volume capacity in Gigabytes.
- **Status**—Volume status.
- **Emulation**—Emulation type for the volume.
- **Host Paths**—Number of masking records for the volume.
- **Reserved**—Indicates whether the volume is reserved.
- **Split**—The name of the associated split.
- **CU Image** —The number of the associated CU image.
- **Base Address**—Base Address.


The following controls are available, however, not all controls may be available for every volume type:

-  —[Viewing CKD volume details](#) on page 204
- **Create**—[Creating volumes](#) on page 178
- **Expand**—[Expanding existing volumes](#) on page 191

-  —[Deleting volumes](#) on page 188
- **Create SG**—[Creating storage groups](#) on page 112
- **Set Volumes > Emulation**—[Setting volume emulation](#) on page 96
- **Set Volumes > Attribute**—[Setting volume attributes](#) on page 195
- **Set Volumes > Identifier**—[Setting volume identifiers](#) on page 196
- **Set Volumes > Status**—[Setting volume status](#) on page 194
- **Set Volumes > Replication QoS**—[QOS for replication](#) on page 197
- **Set Volumes > Set SRDF GCM**—[Setting the SRDF GCM flag](#) on page 434
- **Set Volumes > Reset SRDF/Metro Identity**—[Resetting original device identity](#) on page 432
- **Allocate/Free/Reclaim > Start** —[Managing thin pool allocations](#) on page 244
- **Allocate/Free/Reclaim > Stop**— [Managing thin pool allocations](#) on page 244
- **Configuration > Change Volume Configuration**—[Changing volume configuration](#) on page 190
- **Configuration > Duplicate Volume**—[Duplicating volumes](#) on page 188
- **Configuration > z/OS Map**—[z/OS map from the volume list view](#) on page 333
- **Configuration > z/OS Unmap**—[z/OS unmap from the volume list view](#) on page 334

## Viewing CKD volume details

### Procedure

1. Select the storage system.
2. Select **HOSTS > Mainframe** .
3. Select a CKD volume and click  to open the **Details** view.

---

### Note

Depending on the method you used to open this view, some of the following properties may not appear.

---

The following properties are displayed:

- **Masking Info**—Number of other pools.
- **Storage Groups**—Number of Storage Groups.
- **SRP**—Number of Storage Resource pools (SRPs).
- **CKD Front End Paths**—Number of CKD Front End Paths.
- **RDF Info**—RDF Info.
- **CU Image Number**—CU image number.
- **Split**—Split identifier.

- **Volume 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.
- **Encapsulated Device Flag** — Encapsulated device flag.
- **Encapsulated Device Array** — Encapsulated device array.
- **Encapsulated Device Name** — Encapsulated device name.
- **Status** — Volume status.
- **Reserved** — Whether the volume is reserved.
- **Capacity (GB)** —Volume capacity in GBs.
- **Capacity (MB)** —Volume capacity in MBs.
- **Capacity (CYL)** —Volume capacity in cylinders.
- **Compression Ratio** —Compression ratio.
- **Emulation** — Volume emulation.
- **AS400 Gatekeeper** — AS400 Gatekeeper indication.
- **Symmetrix ID** — Symmetrix system on which the volume resides.
- **Symmetrix Vol ID** — Symmetrix volume name/number.
- **HP Identifier Name** — User-defined volume name (1-128 alpha-numeric characters), applicable to HP-mapped devices. This value is mutually exclusive of the VMS ID.
- **VMS Identifier Name** — Numeric value (not to exceed 32766) with relevance to VMS systems. This value is mutually exclusive of the HP ID.
- **Nice Name** — Nice name generated by Symmetrix Engineuity.
- **WWN** — World Wide Name of the volume.
- **External Identity WWN** — External Identity 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 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).
- **Copy Pace - RDF** — Copy Pace - RDF.
- **Copy Pace - Mirror Copy** — Copy Pace - Mirror Copy.
- **Copy Pace - Clone** — Copy Pace - Clone.
- **Copy Pace - VLUN** — Copy Pace - VLUN.
- **Dynamic Cache Partition Name** — Name of the cache partition.
- **Compressed Size (GB)** — Compressed Size (GB)
- **Compressed Percentage** — Compressed Ratio (%)
- **Compressed Size Per Pool (GB)** — Compressed Size Per Pool (GB)
- **XtremSW Cache Attached** — Indicates whether XtremSW cache is attached to the volume.
- **Base Address** — Base address.
- **AS400 Gatekeeper** — AS400 Gatekeeper indication.
- **Mobility ID Enabled** — Indication if Mobility ID is enabled.
- **GCM** — GCM indication.
- **Optimized Read Miss** — Cacheless read miss status.
- **Persistent Allocation** — Persistent Allocation indication.
- **PowerPath Hosts** — Number of PowerPath hosts.
- **Mounted** — Mounted indication.



- **Process** — Process.
- **Last time used** — Last time used.

The **Details** view links you to views displaying objects contained in and associated with the virtual volume. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Storage Group** opens a view listing the storage groups associated with the volume.

## Viewing CKD volume front end paths

This procedure explains how to view CKD volume front end paths.

### Procedure

1. Select the storage system.
2. Select **Hosts > CU Images** to open the **CU Images** list view.
3. Select the **CU image** and click .
4. Click on the number in the **Number of Volumes** field to open the **CKD Volumes** list view.
5. Select a CKD volume and click  to open its **Details** view.
6. Click on the number in the **CKD Front End Paths** field to open the **CKD Front End Path** list view.
7. The following properties display:
  - **Director Identifier**—Director name.
  - **Port**—Port number.
  - **Base Address**—Assigned base address.
  - **Alias Count**—Number of aliases mapped to the port.
  - **Director Port Status**—Indicates port status.

## Viewing DLDEV volumes

This procedure explains how to view DLDEV volumes.

### Procedure


1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** or **Meta** tab.
3. Filter on DLDEV type.
4. To view the properties and controls, see [Viewing virtual volumes](#) on page 224 or [Viewing meta volumes](#) on page 208.

## Viewing DLDEV volume details

This procedure explains how to view DLDEV volume details.

### Procedure



1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** or **Meta** tab.

3. Filter on DLDEV type.
4. Select a DLDEV volume and click  to open its **Details** view.
5. To view the properties, see [Viewing virtual volume details](#) on page 225 or [Viewing meta volume details](#) on page 209 .

## Viewing masking information

This procedure explains how to view masking information.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** .
3. Select a storage group and click  to open its **Details** view.
4. Click the number next to **Volumes**.
5. Select a volume and click  to open its **Details** view.
6. Click the number next to **Masking Info** to open the volume's **Masking Info** view.

The following properties display:

- **Director Port** — Storage system director and port.
- **Identifier** — Volume identifier name.
- **Type** — Director type.
- **User Generated Name** — User-generated name.
- **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.
- **Port Flag Overrides** — Flag indicating if any port flags are overridden by the initiator: Yes/No.
- **FCID LockDown** — Flag indicating if port lockdown is in effect: Yes/No.
- **Heterogeneous Host** — Whether the host is heterogeneous.
- **LUN Offset** — Whether LUN offset is enabled. This feature allows you to skip over masked holes in an array of volumes.
- **Visibility** — Whether the port is visible to the host.

## Viewing meta volumes

This procedure explains how to view meta volumes.

Meta volumes are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Meta** tab.


Use the this list view to display and manage the volumes. Filter on a volume type.

The following properties display:



- **Name**—Assigned volume name.
- **Type**—Type of volume.
- **Meta Config**—Type of meta volume addressing.
- **Status**—Volume status.
- **Capacity (GB)**—Volume capacity in Gigabytes.
- **Emulation**—Emulation type for the volume.

The following controls are available:

-  —[Viewing meta volume details](#) on page 209
- **Create**—[Creating diskless volumes](#) on page 180
- **Add Member**—[Adding meta members](#) on page 201
- **Dissolve**—[Dissolving meta volumes](#) on page 202
- **Convert**—[Converting meta volumes](#) on page 202
- **Set Volumes > Attribute**—[Setting volume attributes](#) on page 195
- **Set Volumes > Identifier**—[Setting volume identifiers](#) on page 196
- **Set Volumes > Status**—[Setting volume status](#) on page 194
- **Set Volumes > Replication QoS**—[QOS for replication](#) on page 197
- **Configuration > Change Volume Configuration**—[Changing volume configuration](#) on page 190
- **Configuration > Duplicate Volume**—[Duplicating volumes](#) on page 188
- **Configuration > Map**—[Mapping volumes](#) on page 192
- **Configuration > Unmap**—[Unmapping volumes](#) on page 193
- **RecoverPoint > Tag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472
- **RecoverPoint > Untag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472
- **FAST > Pin**—[Pinning and unpinning volumes](#) on page 173
- **FAST > Unpin**—[Pinning and unpinning volumes](#) on page 173
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945
- **Assign Symmetrix Priority**—[Assigning array priority to individual volumes](#) on page 189
- **VLUN Migration**—[VLUN Migration dialog box](#) on page 260
- **Set Optimized Read Miss**—[Setting optimized read miss](#) on page 193


## Viewing meta volume details

This procedure explains how to view meta volume details.

Meta volumes are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.

2. Select **STORAGE > Volumes** and click the **Meta** tab.
3. Select a meta volume and click  to open the **Details** view.

The following properties display:

- **META Members**—Number of Meta members..
- **Storage Groups**—Number of Storage Groups.
- **FBA Front End Paths**—Number of FBA Front End Paths.
- **Back End Paths**—Number of Back End Paths.
- **Volume Name**—Volume name.
- **RDF Info**—RDF Info.
- **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.
- **Encapsulated Device Flag** — Encapsulated device flag.
- **Encapsulated Device Array** — Encapsulated device array.
- **Encapsulated Device Name** — Encapsulated device name.
- **Status** — Volume status.
- **Reserved** — Whether the volume is reserved.
- **Capacity (GB)** —Volume capacity in GBs.
- **Capacity (MB)** —Volume capacity in MBs.
- **Capacity (CYL)** —Volume capacity in cylinders.
- **Emulation** — Volume emulation.
- **Symmetrix ID** — Symmetrix system on which the volume resides.
- **Symmetrix Vol 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 Engenuity.
- **WWN** — World Wide Name of the volume.
- **External Identity WWN** — External Identity World Wide Name of the volume.
- **Mobility ID Enabled** — Indication if Mobility ID is enabled.
- **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** — SRDF 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.
- **GCM** — Indication if GCM is set.
- **SSID** — Subsystem ID.
- **Capacity (Tracks)** — Capacity in tracks.
- **SA Status** — Volume SA status.
- **Host Access Mode** — Host access mode.
- **Pinned** — Whether the volume is pinned.
- **Service State** — Service state.
- **Defined Label Type** — Type of user-defined label.
- **Dynamic RDF Capability** — RDF capability of the volume.
- **Mirror Set Type** — Mirror set for the volume and the volume characteristic of 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.
- **Persistent Allocation** — Persistent Allocation.


There are links to views displaying objects contained in and associated with the volume. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking the number next to **META Members** opens a view listing the members for the meta volume, excluding the meta head.

## Viewing meta volume meta members

This procedure explains how to view meta volume meta members.

Meta volumes are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Meta** tab.
3. Select a meta volume and click  to open its **Details** view.
4. Click the number next to **META Members** to open the **Meta Members** list view.


Use the **Meta Members** list view and manage the members of a meta volume, excluding the meta head.

This list view can be accessed from other volumes that contain meta volumes, that is regular and virtual volumes can contain meta volumes.

The follow properties display:

- **Name**—Meta volume name.
- **Type**—Meta volume configuration.
- **Status**—Volume status.
- **Capacity (GB)**—Volume capacity (GB).

The following controls are available:



-  —[Viewing meta volume member details](#) on page 212
- **Add Meta Member**—[Adding meta members](#) on page 201
- **Remove Meta Member**—[Removing meta members](#) on page 202

## Viewing meta volume member details

This procedure explains how to view meta volume member details.

Meta volumes are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Meta** tab.
3. Select a meta volume and click  to open its **Details** view.
4. Click the number next to **META Members** to open the **Meta Members** list view.
5. Select a meta volume and click  to open its **Details** view.

This list view can be accessed from other volumes that contain meta volumes, that is regular and virtual volumes can contain meta volumes.

Use this view to view meta volume member details.


The following properties display:

- **Physical Name**—Volume's physical name.
- **Volume Identifier**—Volume ID.
- **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.
- **Encapsulated Device Flag** — Encapsulated device flag.
- **Encapsulated Device Array** — Encapsulated device array.
- **Encapsulated Device Name** — Encapsulated device name.
- **Status**—Volume status.
- **Reserved**—Whether the volume is reserved.
- **Capacity (GB)**—Volume capacity in GBs.
- **Capacity (MB)**—Volume capacity in MBs.
- **Capacity (CYL)**—Volume capacity in cylinders.
- **Emulation**—Volume emulation.
- **Stripe Size**—Stripe size.
- **Meta Index**—Meta Index.
- **Symmetrix ID**—Storage system on which the volume resides.
- **Symmetrix Vol 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.
- **External Identity WWN** — External Identity World Wide Name of the volume.
- **Mobility ID Enabled** — Indication that the mobile ID is enabled or not.
- **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.
- **GCM** — GCM indication.
- **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 RecoverPoint is tagged.
- **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.
- **Persistent Allocation** — Persistent Allocation.

## Viewing other pool information

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** tab.
3. Filter on **TDEV** or volume type with **TDEV**, such as **BCV+TDEV**.
4. Select a thin volume and click  to open its **Details** view.
5. Click the number next to **Other Pool Info** to open the **Other Pool Info** view.  
Use this view to view other pool information.

The following properties display:

- **Name**—Thin volume name.
- **Pool Name**—Name of pool.
- **Allocated %**—Percentage of pool allocated to the thin volume.
- **Capacity (GB)**—Amount of pool allocated to the thin volume.

## Viewing private volumes

This procedure explains how to view the properties of private volumes.

### Procedure



1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Private** tab.

Use the this list view to view and manage the volumes. Filter on a volume type.

The following properties display:

- **Name**—Assigned volume name.
- **Type**—Type of volume.
- **Status**—Volume status.
- **Capacity (GB)**—Volume capacity in Gigabytes.
- **Emulation**—Emulation type for the volume.


The following controls are available:

-  — [Viewing private volume details](#) on page 215
- **Create** — [Creating private volumes](#) on page 180
-  — [Deleting volumes](#) on page 188
- **Configuration > Change Volume Configuration**—[Changing volume configuration](#) on page 190
- **Configuration > Duplicate Volume**—[Duplicating volumes](#) on page 188
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945

## Viewing private volume details

This procedure explains how to view private volume details.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Private** tab.
3. Select a private volume and click  to open its **Details** view.

The following properties display:

- **FBA Front End Paths**—Number of FBA Front End Paths.
- **RDF Info**—RDF Info.

- **Volume Name**—Volume name.
- **Back End Paths**—Number of Back End Paths.
- **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.
- **Encapsulated Device Flag** — Encapsulated device flag.
- **Encapsulated Device Array** — Encapsulated device array.
- **Encapsulated Device Name** — Encapsulated device name.
- **Status** — Volume status.
- **Reserved** — Whether the volume is reserved.
- **Capacity (GB)** —Volume capacity in GBs.
- **Capacity (MB)** —Volume capacity in MBs.
- **Capacity (CYL)** —Volume capacity in cylinders.
- **Emulation** — Volume emulation.
- **Symmetrix ID** — Symmetrix system on which the volume resides.
- **Symmetrix VolID** — Symmetrix volume name/number.
- **HP Identifier Name** — User-defined volume name (1-128 alpha-numeric characters), applicable to HP-mapped devices. This value is mutually exclusive of the VMS ID.
- **VMS Identifier Name** — Numeric value (not to exceed 32766) with relevance to VMS systems. This value is mutually exclusive of the HP ID.
- **Nice Name** — Nice name generated by Symmetrix Engineuity.
- **WWN** — World Wide Name of the volume.
- **External Identity WWN** — External Identity 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.
- **GCM** — GCM indication.
- **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 the mirror.
- **Mirror Set DA Status** — Volume status information for each member in the mirror set.
- **Mirror Set Invalid Tracks** — Number of invalid tracks for each mirror in the mirror set.
- **Priority QoS** — Priority value assigned to the volume. Valid values are 1 (highest) through 16 (the lowest).
- **Dynamic Cache Partition Name** — Name of the cache partition.

The **Details** view links you to views displaying objects contained in and associated with the volume. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking the number next to **Back End Paths** opens a view listing the back end paths associated with the volume.

4.

## Viewing regular volumes

This procedure explains how to view regular volumes

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Regular** tab.



Use the this list view to view and manage the volumes. Filter on a volume type.

The following properties display:

- **Name**—Assigned volume name.
- **Type**—Type of volume.
- **Status**—Volume status.
- **Capacity (GB)** — Volume capacity in Gigabytes.

- **Emulation**—Emulation type for the volume.


The following controls are available:

-  —[Viewing regular volume details](#) on page 218
- **Create**—[Creating diskless volumes](#) on page 180
-  —[Deleting volumes](#) on page 188
- **Set Volumes > Attribute**—[Setting volume attributes](#) on page 195
- **Set Volumes > Identifier**—[Setting volume identifiers](#) on page 196
- **Set Volumes > Status**—[Setting volume status](#) on page 194
- **Set Volumes > Replication QoS**—[QOS for replication](#) on page 197
- **Configuration > Change Volume Configuration**—[Changing volume configuration](#) on page 190
- **Configuration > Duplicate Volume**—[Duplicating volumes](#) on page 188
- **Configuration > Map**—[Mapping volumes](#) on page 192
- **Configuration > Unmap**—[Unmapping volumes](#) on page 193
- **Configuration > z/OS Map**—[z/OS map from the volume list view](#) on page 333
- **Configuration > z/OS Unmap**—[z/OS unmap from the volume list view](#) on page 334
- **RecoverPoint > Tag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472
- **RecoverPoint > Untag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472
- **FAST > Pin**—[Pinning and unpinning volumes](#) on page 173
- **FAST > Unpin**—[Pinning and unpinning volumes](#) on page 173
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945
- **Assign Symmetrix Priority**—[Assigning array priority to individual volumes](#) on page 189
- **VLUN Migration**—[VLUN Migration dialog box](#) on page 260
- **Set Optimized Read Miss**—[Setting optimized read miss](#) on page 193

## Viewing regular volume details

This procedure explains how to view regular volume details.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Regular** tab.
3. Select a regular volume and click  to open the **Details** view.

The **Details** view allows you to view and manage a volume.

### Properties panel

The following properties display:

- **Masking Info**—Masking Info.
- **Storage Groups**—Number of Storage Groups.
- **FBA Front End Paths**—Number of FBA Front End Paths.
- **RDF Info**—RDF Info.
- **Volume Name**—Volume name.
- **Back End Paths**—Number of Back End Paths.
- **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.
- **Encapsulated Device Flag** — Encapsulated device flag.
- **Encapsulated Device Array** — Encapsulated device array.
- **Encapsulated Device Name** — Encapsulated device name.
- **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.
- **AS400 Gatekeeper** — AS400 Gatekeeper indication.
- **Symmetrix ID** — Symmetrix system on which the volume resides.
- **Symmetrix Volume ID** — Symmetrix volume name/number.
- **HP Identifier Name** — User-defined volume name (1-128 alpha-numeric characters), applicable to HP-mapped devices. This value is mutually exclusive of the VMS ID.
- **VMS Identifier Name** — Numeric value (not to exceed 32766) with relevance to VMS systems. This value is mutually exclusive of the HP ID.
- **Nice Name** — Nice name generated by Symmetrix Engineuity.
- **WWN** — World Wide Name of the volume.
- **External Identity WWN** — External Identity 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.
- **GCM** — GCM indication.
- **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 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.
- **Copy Pace - RDF** — Copy pace priority during RDF operations.
- **Copy Pace - Mirror Copy** — Copy pace priority during mirror operations.
- **Copy Pace - Clone** — Copy pace priority during clone operations.
- **Copy Pace - VLUN** — Copy pace priority during virtual LUN operations.
- **XtremSW Cache Attached** — Indicates whether XtremSW cache is attached to the volume.
- **Optimized Read Miss** — Cacheless read miss status.
- **Persistent Allocation** — Persistent Allocation indication.

The **Details** view links you to views displaying objects contained in and associated with the volume. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking the number

next to **Storage Groups** opens a view listing the storage groups associated with the volume.

## Viewing reserved volumes

### Procedure

1. Select the storage system.
2. In the dashboard, click the **System Health** tab.
3. In the **Action** panel, click **View Reservations**.
4. Select the reservation and click .
5. Click the number next to **Reserved Volumes**.

The **Reserved Volume** list view is displayed.

Use the **Reserved Volumes** list view to display and manage the volumes held in a reservation.

The following properties display:


- **Name**—Assigned volume name.
- **Type**—Type of volume.
- **Capacity (GB)**—Volume capacity in Gigabytes.
- **Status**—Volume status.
- **Reserved**—Indicates whether the volume is reserved.
- **Emulation**—Emulation type for the volume.

The following controls are available:

-  — [Viewing reserved volume details](#) on page 221.

## Viewing reserved volume details

### Procedure

1. Select the storage system.
2. In the dashboard, click the **System Health** tab.
3. In the **Action** panel, click **View Reservations**.
4. Select the reservation and click .  
Opens the **Details** view.
5. Click the number next to **Reserved Volumes** to open the **Reserved Volumes** list view .

The following properties display:

- **Name**—Volume name.
- **Volume Identifier** —Volume identifier.
- **Type**—Volume configuration.

- **Encapsulated Volume**—Whether external volume is encapsulated. 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**—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 consistency 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.

## Viewing SAVE volumes

This procedure explains how to view SAVE volumes.


### Procedure

1. Select **STORAGE > Volumes** and click the **Private** tab.
2. Filter on SAVE type.
3. To view the properties and controls, see [Viewing private volumes](#) on page 215.

## Viewing SAVE volume details

This procedure explains how to view SAVE volume details.

### Procedure


1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Private** tab.
3. Filter on SAVE type.
4. Select a SAVE volume and click  to open its **Details** view.

5. To view the properties, see [Viewing private volume details](#) on page 215.

## Viewing storage resource pool information

This procedure explains how to view storage resource pool information.

### Procedure

1. Select the storage system.
2. Select **Storage > Volumes** to open the **Volumes** list view.
3. Select the volume and click  to open its **Details** view.
4. Click the number next to **SRP** to go to the **Storage Resource Pool** view for the volume.

The following properties display:

- **Name**—Volume name.
- **SRP Name**—Storage resource pool name.
- **Allocated**—Volume capacity allocated.
- **Capacity**—Total volume capacity.
- **Allocated %**—Percent of volume used.

## Viewing thin volumes

This procedure explains how to view thin volumes.


### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** tab.
3. Filter on a thin volume type, such as TDEV.
4. To view the properties and controls, see [Viewing virtual volumes](#) on page 224.

## Viewing thin volume details

This procedure explains how to view thin volume details.

### Procedure


1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** tab.
3. Filter on a thin volume type, such as TDEV.
4. Select a thin volume and click  to open its **Details** view.
5. To view the properties, see [Viewing virtual volume details](#) on page 225.

## Viewing thin volume bound pool information

This procedure explains how to view thin volume bound pool information.

### Procedure

1. Select the storage system.

2. Select **STORAGE > Volumes** and click one of the panels.
3. Select the thin volume and click  to open its **Details** view.
4. Click the number next to **Bound Pool Info** to open the **Bound Pool Info** view.

The following properties display:

- **Name**—Thin volume name.
- **Pool Name**—Name of pool.
- **Allocated %**—Percentage of pool allocated to the thin volume.
- **Capacity (GB)**—Capacity in GB.
- **Allocated (GB)**—Number of GB allocated from the pool for exclusive use by the thin volume.
- **Subscription %**—Ratio between the DATA volume pool's enabled capacity and the thin volume subscribed capacity.
- **Written (GB)**—Number of allocated GB in the DATA volume pool that the thin volume has used.
- **Shared Tracks**—Whether tracks are shared between thin volumes.
- **Persistent Allocation**—Indicates persistent allocations: All, some, or none.

## Viewing virtual volumes

This procedure explains how to view virtual volumes.

### Procedure



1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** tab.

Use the this list view to view and manage the volumes. Filter on a volume type.

The following properties display:

- **Name**—Assigned volume name.
- **Type**—Type of volume.
- **Emulation**—Emulation type for the volume.
- **Capacity (GB)**—Volume capacity in Gigabytes.
- **Status**—Volume status.

The following controls are available:

-  —[Viewing virtual volume details](#) on page 225
- **Create** —[Creating VDEV volumes](#) on page 187
-  —[Deleting volumes](#) on page 188
- **Set Volumes > Attribute**—[Setting volume attributes](#) on page 195
- **Set Volumes > Identifier**—[Setting volume identifiers](#) on page 196
- **Set Volumes > Status**—[Setting volume status](#) on page 194
- **Set Volumes > Replication QoS**—[QoS for replication](#) on page 197




- **Configuration > Change Volume Configuration**—[Changing volume configuration](#) on page 190
- **Configuration > Duplicate Volume**—[Duplicating volumes](#) on page 188
- **Configuration > Map**—[Mapping volumes](#) on page 192
- **Configuration > Unmap**—[Unmapping volumes](#) on page 193
- **RecoverPoint > Tag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472
- **RecoverPoint > Untag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472
- **Allocate/Free/Reclaim > Start** —[Managing thin pool allocations](#) on page 244
- **Allocate/Free/Reclaim > Stop**—[Managing thin pool allocations](#) on page 244
- **FAST > Bind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257
- **FAST > Unbind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257
- **FAST > Rebind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257
- **FAST > Pin**—[Pinning and unpinning volumes](#) on page 173
- **FAST > Unpin**—[Pinning and unpinning volumes](#) on page 173
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945
- **VLUN Migration**—[VLUN Migration dialog box](#) on page 260
- **Set Optimized Read Miss**—[Setting optimized read miss](#) on page 193

## Viewing virtual volume details

This procedure explains how to view virtual volume details.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Virtual** panel.
3. Select a volume and click  to open its **Details** view.

The **Details** view allows you to view and manage a volume.

The following properties display:

- **Bound Pool Info**—Number of bound pools.
- **Other Pool Info**—Number of other pools.
- **Masking Info**—Number of other pools.
- **Storage Groups**—Number of Storage Groups.
- **FBA Front End Paths**—Number of FBA Front End Paths.
- **RDF Info**—RDF Info.
- **Volume 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.
- **Encapsulated Device Flag** — Encapsulated device flag.
- **Encapsulated Device Array** — Encapsulated device array.
- **Encapsulated Device Name** — Encapsulated device name.
- **Status** — Volume status.
- **Reserved** — Whether the volume is reserved.
- **Capacity (GB)** — Volume capacity in GBs.
- **Capacity (MB)** — Volume capacity in MBs.
- **Capacity (CYL)** — Volume capacity in cylinders.
- **Emulation** — Volume emulation.
- **AS400 Gatekeeper** — AS400 Gatekeeper indication.
- **Symmetrix ID** — Symmetrix system on which the volume resides.
- **Symmetrix Vol 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 Engenuity.
- **WWN** — World Wide Name of the volume.
- **External Identity WWN** — External Identity 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.
- **GCM** — GCM indication.
- **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 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.
- **Compressed Size (GB)** — Compressed Size (GB)
- **Compressed Ratio (%)** — Compressed Ratio (%)
- **Compressed Size Per Pool (GB)** — Compressed Size Per Pool (GB)
- **XtremSW Cache Attached** — Indicates whether XtremSW cache is attached to the volume.
- **Optimized Read Miss** — Cacheless read miss status.
- **Persistent Allocation** — Persistent Allocation indication.

The **Details** view links you to views displaying objects contained in and associated with the virtual volume. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking **Storage Group** opens a view listing the two storage groups associated with the volume.

## Viewing volume back end paths

This procedure explains how to view volume back end paths.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click one of the panels.
3. Select a volume and click  to open its Details view.
4. Click the number next to **Back End Paths**.

This view allows you to view the back end paths associated with the volume.


The following properties display:

- **Name**—Name.
- **Director Identifier**—Director identifier.
- **DA Interface**—DA interface ID.
- **SCSI ID**—Disk SCSI ID.
- **DA Volume Number**—DA volume ID.
- **Hyper Number**—Hyper ID.
- **Hyper Capacity** —Hyper capacity.
- **Member Status**—Hyper member status.
- **Member Number**—Hyper member number.
- **Disk Group Pretty Name**—Name of disk group.
- **Disk Capacity** —Capacity of disk.
- **Spindle**—Spindle ID.

## Viewing volume FBA front end paths

This procedure explains how to view volume FBA front end paths.

### Procedure

1. Select the storage system.
2. Select a volume and click  to open its **Details** view.
3. Click the number next to **FBA Front End Paths** to open the **FBA Front End Paths** list view.

Use the **FBA Front End Paths** list view to view the FBA front end paths associated with a volume.

The following properties display:


- **Director Identifier**—Director name.
- **Port**—Port number.
- **VBus**—VBus number.
- **TID**—Disk SCSI ID.
- **Symm LUN**—Symmetrix LUN number.
- **PDeVName**—Physical device name.
- **Director Port Status** — Director port status.

## Viewing volume RDF information

This procedure explains how to view volume RDF information.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click one of the tabs.

3. Select a volume and click  to open its **Details** view.
4. Click the number next to **RDF Info** to open the **RDF Info** list view.

The following properties display:

- **Remote SymmID**—Remote Symmetrix serial ID.
- **RDev**—Symmetrix volume name.
- **RDev Config**—Volume configuration.
- **Capacity (GB)**—Volume capacity.
- **RDFG**—RDF group containing the volume.
- **Pair State**—State of the pair of which the volume is part.
- **RDF Feature**—SRDF copy type.
- **CSRMT** —CSRMT — 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

- **R1 Inv**—Number of invalid tracks on the R1 volume.
- **R2 Inv**—Number of invalid tracks on the R2 volume.
- **RA Status**—Status of the RDF director.
- **Link Status**—Indicates link state.
- **RDF State**—Volume RDF state.
- **Remote RDF State**—Remote volume RDF state.
- **RDF Status**—Volume RDF status.
- **Device Config RDFA WPACE Exempt**—Indicates is write pacing exemption capability is enabled or disabled.
- **Effective RDFA WPACE Exempt**— indicates if effective write pacing exemption capability is enabled or disabled.

## Select Volume Range dialog box

Use this dialog box to select the range of volumes for the operation.

The following properties display:

**Volume Range**—Range of volumes.

**CU Image Number**—CU image containing the volumes.

**SSID**—Subsystem ID assigned to the volumes.

**Base Address**—Base addresses assigned to the volumes.

**Aliases**—Aliases assigned to the volumes.

## Advanced Options dialog

Refer to the parent help topic for information on the **Advanced Options** dialog.

## Viewing disk groups

### Procedure



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 ; format is : `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:

-  — [Viewing disk group details](#) on page 230
- **Rename**—Change the name of a disk group.
-  — [Deleting disk groups](#) on page 237

## Viewing disk group details


### Procedure

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, click , and do one of the following:
  - Click the **Details** tab.  
The following properties display:
    - **Name**—Name of disk group ; format is : `number -- name`.
    - **Technology**—Technology type for the disk group.
    - **Used Capacity (GB)**—Used capacity.
    - **Free Capacity (GB)**—Free capacity.
    - **Total Capacity (GB)**—Total capacity.
    - **Speed (RPM)**—Speed of the disks in the group.
    - **Form Factor**—Form factor.

- **Disk Location**—Whether the disks in the group are internal or external.
- **Number of Disks**—Number of disks.
- **Number of Spare Disks**—Number of spare disks.
- Click the **Disk Group Usage Report** tab.  
A visual representation of used capacity as a percentage of overall capacity is displayed.

## Viewing disks in disk group


### Procedure

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, click , and click the **Details** tab.
4. Click the number next to **Number of 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:

-  — [Viewing disk details](#) on page 231
- **Remove Disk**—[Removing disks from disk groups](#) on page 236

## Viewing disk details

### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group, click , and click the **Details** tab to open its **Details** view.
4. Click the number next to **Number of Disks** to open the **Disks for Disk Group** list view.
5. Select the disk group from the list, click , and do one of the following:


- Click the **Details** tab.  
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 **Details** 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 **Number of Hypers** opens the view listing the hypers contained in the disk.


- Click the **Disk Group Usage Report** tab.  
A visual representation of used capacity as a percentage of overall capacity is displayed.

## Viewing disk hyper volumes

### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group, click , and click the **Details** tab.



4. Click the number next to **Number of Disks** to open the **Disks for Disk Group** list view.
5. Select a disk, click , and click the **Details** tab.
6. Click the number next to **Number of Hypers** to open the **Hypers for Disk** list view.

Use the **Hypers for Disk** list view to view the hyper volumes in a disk group. The following properties display:



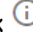
- **Hyper**—Volume hyper number.
- **Volumes**—Disk adapter logical volume number (1 - n).
- **Hyper Type**—Hyper type.
- **Mirror**—Mirror position of hyper.
- **Capacity (GB/Cyl)**—Disk capacity in GB/Cylinders.
- **Symm Volume**—Symmetrix volume number.
- **Hyper Status**—Hyper status.
- **Emulation**—Emulation of hyper volume.

The following control is available:

-  —[Viewing hyper volume details](#) on page 233

## Viewing hyper volume details

### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group, click , and click the **Details** tab.
4. Click the number next to **Number of Disks** to open the **Disks for Disk Group** list view.
5. Select a disk, click , and click the **Details** tab.
6. Click the number next to **Number of Hypers** to open the **Hypers for Disk** list view.
7. Select a hyper volume and click  to open its **Details** view.
8. Use the hyper volume **Details** view to view the properties of a hyper volume.

The following properties display:

- **Hyper Number**—Volume hyper number.
- **DA Volume**—Disk adapter logical volume number (1 - n).
- **Hyper Type**—Hyper type.
- **Mirror**—Mirror position of hyper.
- **Capacity (GB/Cyl)**—Disk capacity in GB/Cylinders.
- **Symm Volume**—Symmetrix volume number.

- **Raid Group**—RAID-S group number.
- **Original Mirror**—Mirror position of hyper.
- **Hyper Status**—Hyper status.
- **Emulation**—Emulation of hyper volume.



## Viewing list for a hyper type

Depending on your selection, a list is displayed for one of the following :

- MetaHypers
- Raid5Hyper
- Raid5MetaHyper
- Raid6Hyper
- Raid6MetaHyper
- Hypers
- TWM
- MetaTWM
- MetaMembers



## Viewing volumes for disk

### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group, click , and click the **Details** tab.
4. Click the number next to **Number of Disks** to open the **Disks for Disk Group** list view.
5. Select a disk and click  to open the details view for the disk.
6. Click the number next to **Number of Volumes** to open the **Volumes** view.

## Viewing paths for disks

### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group, click , and click the **Details** tab.
4. Click the number next to **Disks** to open the **Disks for Disk Group** list view.
5. Select a disk, click , and click the **Details** tab.
6. Click the number next to **Number of Paths** to open the **Paths for Disk** list view.


Use the **Paths for Disk** list view to view the paths for a disk.

The following properties display:

- **Dir**—Director Identifier. Possible values are a director number or the word "Multi," which indicates that the hyper can see multiple directors.
- **Port**—Director port number.
- **Remote Port WWN**—World Wide Name of the port.
- **Active Path**—Whether active path is being used (True/False).
- **Failover**—Whether failover is being used (True/False).

## Viewing spare disks in disk group

### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group and click .
4. In the **Details** panel, click the number next to **Number of Spare Disks** to open the **Spare Disks for Disk Group** view.

Use the **Spare Disks for Disk Group** view to view the spare disks in a disk group.

The following properties display:


- **Dir**—Director ID.
- **Int**—DA SCSI path.
- **TID**—Disk SCSI ID.
- **Hypers**—Number of hypers.
- **Disk Group**—Disk group number where disk is contained.
- **Speed (RPM)**—Physical disk revolutions per minute.
- **Total Capacity (GB)**—Total disk in GB.
- **Failed Dir**—Failed disk director ID.
- **Failed DA Number**—Failed disk DA number.
- **Failed DA Int**—Failed disk DA SCSI path.
- **Failed Disk SCSI ID**—Failed disk SCSI ID.
- **Failed Spindle ID**—Failed disk Spindle ID.


The following controls are available:

-  — [Viewing spare disk details](#) on page 235

## Viewing spare disk details

### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups** to open the **Disk Groups** list view.
3. Select the disk group and click .

4. In the **Details** panel, click the number next to **Number of Spare Disks** to open the **Spare Disks for Disk Group** view.
5. Select a disk and click  to open the its **Details** view.

Use the disk **Details** view to view the properties of a spare disk.  
The following properties display:

- **Spindle**—Spindle ID
- **Dir**—Director ID
- **Int**—DA SCSI path
- **TID**—Disk SCSI ID
- **External WWN**—External world wide name
- **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
- **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 Block**—Actual number of disk blocks
- **Block Size**—Size of each disk block
- **Total Capacity (GB)**—Total disk capacity in Gigabytes
- **Free Capacity (GB)**—Free disk capacity in Gigabytes
- **Actual Capacity (GB)**—Actual disk capacity in Gigabytes
- **Used Capacity (GB)**—Used disk capacity in Gigabytes
- **Used (%)**—Percentage of used disk capacity to the total disk capacity
- **Rated Disk Capacity (GB)**—Rated disk capacity in Gigabytes
- **Spare Disk**—Indication if the disk is a spare disk
- **Encapsulated**—Indication if the disk is encapsulated
- **Disk Service State**—Disk service state

## Removing disks from disk groups

---


### Note

Only empty external disk groups can be deleted.

---

### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups**.


3. Select the disk group from the list and click  to open its **Details** view.
4. From the **Details** panel, click the number next to **Number of Disks** to open the **Disks** view.
5. Select a disk from the list and click **Remove Disk**.
6. Click **OK**.

## Deleting disk groups

### Before you begin

Only empty external disk groups can be deleted.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups** to open the **Disk Groups** list view.
3. Select a disk group and click .
4. Click **OK**.

## Renaming disk groups


### Procedure

1. Select the storage system.
2. Select **STORAGE > Disk Groups**.
3. Select the disk group and click **Rename**.
4. Type the new disk group name and click **OK**.

## Creating DATA volumes

This procedure explains how to create DATA volumes on storage systems running Enginuity version 5876.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click  to open its **Details** view.
4. Click the number next to **Number of Data Volumes**.
5. Click **Create Volumes**.
6. Select **DATA** as the **Configuration**.
7. Select the **Disk Technology**.  
**External** disk technology is an option if the storage system has FTS (Federated Tiered Storage) enabled and available external storage.
8. Select the **Emulation** type.
9. Select the RAID **Protection** level.
10. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**. You can also manually enter a volume capacity.

11. To add the new volumes to a specific thin pool, select one from **Add to Pool**. Pools listed are filtered on technology, emulation, and protection type.
12. Click **Advanced Options** to continue setting the advanced options, as described next.  
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 the 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**.
  - d. Click **APPLY**.
13. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.





## Activating and deactivating DATA volumes

### Before you begin

You can only activate deactivated DATA volumes with used tracks.

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.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Click the **Private** tab.
4. Filter on DATA type.
5. Do one of the following:
  -  Select one or more volumes, click , and select **Set Volumes > Activate**.
  -  Select one or more volumes, click , and select **Set Volumes > Deactivate**.
6. Click **OK**.

## Enabling and disabling DATA volumes



### Before you begin

To disable a volume, all sessions must be terminated, and have no used tracks.

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.



#### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Click the **Private** tab.
4. Filter on DATA type.
5. Do one of the following:
  - Select one or more volumes, click , and select **Set Volumes > Enable**.
  - Select one or more volumes, click , and select **Set Volumes > Disable**.
6. Click **OK**.

## Start and stop draining DATA volumes

This procedure explains how to start or stop draining DATA volumes.

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Click the **Private** tab.
4. Filter on DATA type.
5. Do one of the following:
  - Select one or more volumes, click , and select **Draining > Start**.
  - Select one or more volumes, click , and select **Draining > Stop**.
6. Click **OK**.

## Viewing DATA volumes

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Private** tab.
3. Filter on DATA type.
4. To view the properties and controls, see [Viewing private volumes](#) on page 215.

## Viewing DATA volume details

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click the **Private** tab.

3. Filter on DATA type.
4. Select a DATA volume and click  to open its **Details** view.
5. To view the properties, see [Viewing private volume details](#) on page 215.

## Creating thin pools

When creating thin pools, Unisphere 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, 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:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3. Click **Create** to open the **Create Thin Pool** dialog box.

When this dialog box first opens, the chart displays the configured and unconfigured space on the selected storage 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 **Thin Pool Name**.

Thin pool names can contain up to 12 alpha-numeric characters. The only special characters allowed are the hyphen (-) and the underscore ( \_ ); however, the name cannot start or end with a or hyphen or 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 number of volumes, volume capacity and capacity unit.
9. Click **Advanced Options** - see [Creating or Expanding or Modifying thin pools](#) on page 255.
10. Verify your selections in the **Create Thin Pool - Summary** page, 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 920 and [Previewing jobs](#) on page 920.



- Expand **Add to Job List** and click **Run Now** to create the pool now.

## Expanding thin pools

### Before you begin

Unisphere supports best practices, which state that volumes from different drive technologies should not be mixed in the same thin pool. To this end, Unisphere 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.

Expanding thin pools refers to the process of increasing the amount of pool storage accessible to a thin volume by either adding a predefined capacity to the pool, or by increasing the pool's capacity by a percentage.

To expand a thin pool:

### Procedure

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. Click **Advanced Options** - see [Creating or Expanding or Modifying thin pools](#) on page 255.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List** and click **Run Now** to create the pool 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 238.
- The drain must not cause the enabled volumes to end up with greater than 90% utilization in the pool. To calculate this, add the total used tracks on the enabled volumes and the total used tracks on the volumes that will be drained and divides this sum by the total number of tracks on all the enabled volumes. If the result is greater than 90% the drain request is blocked.

- The number of volumes that are draining at any time are limited to 20% total of the number of volumes to drain (or draining) plus the number of enabled volumes. This limits the impact on the system.
- This feature is only supported on storage systems running Enginuity 5876 or higher.

To drain thin pools:

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click  to open its **Details** view.
4. Click the number next to **Number of Data Volumes** to open the **DATA Volumes** list view.
5.  Select one or more volumes, click , and select **Start Draining**.
6. Click **OK**.  
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 , 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 5 to 9 until all the volumes are drained to an acceptable percentage.

## Starting and stopping thin pool write balancing



#### 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 5876 or higher.

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.

To write balance a thin pool:

#### Procedure


1. Select the storage system
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3.  Start write balancing by clicking , and clicking **Start Write Balancing**.

4. Click **OK**.

This will put the pool in a Balancing state.

5. Monitor the balancing until it reaches an acceptable percentage

- 6.

Select the thin pool, click , and select **Stop Write Balancing**.

7. Click **OK**.

## Deleting thin pools

Before you begin:

You can only delete empty thin pools. For instructions, refer to [Adding or removing thin pool members](#) on page 243.

To delete a thin pool:

### Procedure

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 **OK**.

## Adding or removing thin pool members


This procedure explains how to add or remove members from a thin pool.

Before you begin:

- The storage system must be running Enginuity 5876.
- Before you can remove a thin pool member (data volume), you must first [disable](#) it.
- Unisphere 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 or remove thin pool members:

### Procedure

1. Select the storage system.
2. Select **Storage > Thin Pools** to open the Thin Pools list view.
3. Select the thin pool and click  to open the thin pool's details view.
4. Click the number next to **Number of Data Volumes** to open the **DataVolumes** view.
5. Click **Add Volumes to Pool** to open the **Add Volumes to Thin Pool** wizard.
  - a. 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.
- **Volume Configuration** — Filters the list for the specified configuration.
- **Emulation** — Filters the list for the specified emulation.

b. Click **NEXT**.

c. In the **Available Volumes** table, select the volumes.

d. Deselect one or more the previously selected volumes to remove a volume.

e. Click **OK**.

## Enabling and disabling thin pool members

### Procedure

1. Select the storage system.
2. Select **Storage > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click **SRDF Groups** to open the thin pool's details view.
4. Click the number next to **DATA Volumes**.
5. Do one of the following:
  - To enable members, select them and click **Enable**.
  - To disable members, select them and click **Disable**.
6. Click **OK**.

## Managing thin pool allocations

### Before you begin

- You can only allocate thin pool capacity to bound thin volumes.
- 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.

The following describes how to start and stop allocating thin pool capacity from the Volumes view.



### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Select the volume type by selecting a tab.
4. Do one of the following:

- To start thin pool allocation:

■

Select one or more volumes, click , and select **Allocate/Free/Reclaim > Start**.

- Select **Allocate Volumes**, **Free Volumes**, or **Reclaim Volumes**. If you select **Allocate Volumes**, 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. If you select **Reclaim Volumes**, you can optionally specify to reclaim persistent capacity by selecting the **Reclaim persistent capacity** option.
  - To stop thin pool allocation:
    -  Select one or more volumes, click more , and select **Allocate/Free/Reclaim > Stop**.
    - Select **Stop Allocate**.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing thin pools


### Procedure

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3. The **Thin Pools** list view allows you to view and manage thin pools on a storage 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:

-  —[Viewing thin pool details](#) on page 246
- **Create** —[Creating thin pools](#) on page 240
- **Modify**—[Creating or Expanding or Modifying thin pools](#) on page 255
- **Expand** —[Expanding thin pools](#) on page 241
- **Delete**—[Deleting thin pools](#) on page 243
- **Start Write Balancing**—[Starting and stopping thin pool write balancing](#) on page 242
- **Stop Write Balancing**—[Starting and stopping thin pool write balancing](#) on page 242

- **Bind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257

## Viewing thin pool details

### Procedure

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3. Select the pool and click  to open its **Pool Usage Report** or **Details** panel.

The **Pool Usage Report** panel provides a graphic representation of the thin pool's allocation as a percentage.

The following properties display in the **Details** panel:


- **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.
- **Emulation** — Emulation type for the pool.
- **Total Capacity (GB)** — Total capacity of the pool.
- **Free Capacity (GB)** — Free capacity in the pool.
- **Enabled Capacity (GB)** — Sum of capacity of all enabled DATA volumes in the pool.
- **Allocated Capacity (GB)** — Pool capacity allocated to thin volumes.
- **Allocated %** — Percent of pool used.
- **Maximum Subscription Set** — Enable oversubscription for the pool.
- **Maximum Subscription** — Acceptable oversubscription ratio for the pool.
- **Subscription %** — Current subscription percentage.
- **State** — Pool state (Enabled, Disable, Balancing).
- **Rebalance Variance** — Target volume utilization variance for the rebalancing algorithm. The rebalancing algorithm attempts to level data distribution in a pool so that the percentage utilization of any volume in the pool is within the target variance of the percentage utilization of any other volume in the pool. Possible values range from 1 to 50%, with the default value being 1%. This field is only available when creating a thin pool on a Symmetrix system running Enginuity 5876 or higher.
- **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 5876 or higher.
- **Pool Capacity Reserved** — 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.
- **Number of Bound Volumes** — Number of thin volumes bound to the pool.
- **Number of Data Volumes** — Number of data volumes bound to the pool.
- **Number of Enabled Volumes** — Number of enabled DATA volumes in the pool.
- **Number of Disabled Volumes** — Number of disabled DATA volumes in the pool.
- **Disk Location** — Whether the disk group is internal to the storage system or an external storage system or storage device.

You can view 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 the number next to **Number of Data Volumes** opens a view listing the DATA volumes in the pool.

## Viewing bound volumes for a thin pool

### Procedure

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click  to open its **Details** view.
4. Click the number next to **Number of Bound Volumes**.

The following properties display:

- **Name**—Assigned volume name.
- **Emulation**—Emulation type for the volume.
- **Configuration**—Volume configuration.
- **Capacity (GB)**—Volume capacity in Gigabytes.
- **Allocated (GB)**—Number of GBs from the pool allocated for exclusive use by the volume.
- **Written (GB)**—Number of allocated GBs in the pool that the thin volume has actually used.
- **Shared Tracks**—Whether the volume shares tracks with other thin volumes.

The following controls display:

- **Create Volumes**—[Creating thin volumes](#) on page 184
- **Bind** —[Binding/Unbinding/Rebinding thin volumes](#) on page 257
- **Unbind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257
- **Configuration > Change Volume Configuration**—[Changing volume configuration](#) on page 190

- **Configuration > Map**—[Mapping volume operations](#) on page 97
- **Configuration > Unmap**—[Mapping volume operations](#) on page 97
- **Configuration > z/OS Map**—[Mapping volume operations](#) on page 97
- **Configuration > z/OS Unmap**—[Mapping volume operations](#) on page 97
- **Set Volume > Emulation** —[Setting volume emulation](#) on page 96
- **Set Volume > Attributes**—[Setting volume attributes](#) on page 195
- **Set Volume > Identifiers**—[Setting volume identifiers](#) on page 196
- **Set Volume > Status**—[Setting volume status](#) on page 194
- **FAST > Rebind**—[Binding/Unbinding/Rebinding thin volumes](#) on page 257
- **FAST > Pin**—[Pinning and unpinning volumes](#) on page 173
- **FAST > Unpin**—[Pinning and unpinning volumes](#) on page 173
- **Allocate/Free/Reclaim > Start**—[Managing thin pool allocations](#) on page 244
- **Allocate/Free/Reclaim > Stop** —[Managing thin pool allocations](#) on page 244
- **RecoverPoint > Tag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472
- **RecoverPoint > Untag**—[Tagging and untagging volumes for RecoverPoint \(storage group level\)](#) on page 472
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945
- **VLUN Migration**—[VLUN Migration dialog box](#) on page 260

5. Select a bound volume and click  to open its **Details** view.

The following properties display:

- **Name**—Volume name.
- **Physical name**—Physical name.
- **Volume Identifier**—Volume Identifier.
- **Type**—Volume configuration.
- **Encapsulated Volume**—Indication whether the volume is encapsulated.
- **Encapsulated WWN**—Encapsulated World Wide Name.
- **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 volumes. 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 Engenuity.
- **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 CG 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)**—Capacity.
- **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**—Indication whether 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.

- **Optimized Read Miss**—Optimized Read Miss.
- **Compressed Size (GB)**—Compressed Size.
- **Compressed Percentage**—Compressed percentage.
- **Compressed Size Per Pool**—Compressed size per pool.

## Viewing DATA volumes for a thin pool


### Procedure

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.
3. Select the thin pool and click  to open its **Details** view.
4. Click the number next to **Number of Data Volumes**.

The following properties display:

- **Name**—Name of the DATA volume.
- **Emulation**—Volume emulation.
- **Configuration**—Volume configuration.
- **Used (%)**—Percent of the volume used.
- **Used (GB)**—Space used.
- **Free**—Free space on the volume.
- **Status**—Volume status.
- **Session Status**—Session status (Active, or Inactive).



The following controls are available:

-  —[Viewing details on DATA volumes in thin pools](#) on page 250
- **Create Volumes**—[Creating DATA volumes](#) on page 179
- **Add Volumes to Pool**—[Adding or removing thin pool members](#) on page 243
- **Remove**—[Adding or removing thin pool members](#) on page 243
- **Enable**—[Enabling and disabling thin pool members](#) on page 244
- **Disable**—[Enabling and disabling thin pool members](#) on page 244
- **Activate**—[Activating and deactivating DATA volumes](#) on page 238
- **Deactivate**—[Activating and deactivating DATA volumes](#) on page 238
- **Start Draining**—[Draining thin pools](#) on page 241
- **Stop Draining**—[Draining thin pools](#) on page 241

## Viewing details on DATA volumes in thin pools

### Procedure

1. Select the storage system.
2. Select **STORAGE > Thin Pools** to open the **Thin Pools** list view.

3. Select the thin pool and click  to open its **Details** view.
4. Click the number next to **Number of Data Volumes**.
5. Select a DATA volume and click  to open its **Details** view.


The following properties display:

- **Name**—Volume name.
- **Type**—Volume configuration.
- **Encapsulated Volume**—Indication whether the volume is encapsulated.
- **Encapsulated WWN**—Encapsulated World Wide Name.
- **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 volumes. 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 Engenuity.
- **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 CG 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)**—Capacity.
- **SSID**—Subsystem ID.

- **Capacity (Tracks)**—Capacity in tracks.
- **SA Status**—Volume SA status.
- **Host Access Mode**—Host access mode.
- **Pinned**—Whether the volume is pinned.
- **Service State**—Service state.
- **Defined Label Type**—Type of user-defined label.
- **Dynamic RDF Capability**—RDF capability of the volume.
- **Mirror Set Type**—Mirror set for the volume and the volume characteristic of 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.
- **XtremSWCache Attached**—Whether volume is attached to XtremSW cache.
- **Compression Delta (GB)**—Difference between volume allocation and uncompressed data.

## Viewing other volumes for thin pools


### Procedure

1. Select the storage system.
2. Select **Storage > Thin Pools**.
3. Select the pool and click .
4. Click the number next to **Other Volumes**.
5. Use the **Other Volumes for Thin Pool** list view to display and manage other volumes bound to a thin pool.

The following properties display:

- **Name** — Assigned volume name.
- **Pool Name** — Pool to which the volume is bound.
- **% Allocated** — Percentage of space allocated in the pool.
- **Allocated Capacity** — Amount of space allocated in the pool.

The following controls are available:

-  — [Viewing thin volume details](#) on page 223
- **Create** — [Creating thin volumes](#) on page 184
- **Bind** — [Binding/Unbinding/Rebinding thin volumes](#) on page 257
- **Unbind** — [Binding/Unbinding/Rebinding thin volumes](#) on page 257

- **Untag for RecoverPoint** — [Tagging and untagging volumes for RecoverPoint \(volume level\)](#) on page 472
- **Tag for RecoverPoint** — [Tagging and untagging volumes for RecoverPoint \(volume level\)](#) on page 472
- **Unpin** — [Pinning and unpinning volumes](#) on page 173
- **Pin** — [Pinning and unpinning volumes](#) on page 173
- **Assign Symmetrix Priority** — [Assigning array priority to individual volumes](#) on page 189
- **Unmap** — [Unmapping volumes](#) on page 193
- **Map** — [Mapping volumes](#) on page 192
- **Assign Dynamic Cache Partition** — [Assigning dynamic cache partitions](#) on page 945
- **Stop Allocate/Free/Reclaim** — [Managing thin pool allocations](#) on page 244
- **Set Volume Status** — [Setting volume status](#) on page 194
- **Set Volume Identifiers** — [Setting volume identifiers](#) on page 196
- **Set Volume Attributes** — [Setting volume attributes](#) on page 195
- **Change Volume Configuration** — [Changing volume configuration](#) on page 190
- **Rebind** — [Binding/Unbinding/Rebinding thin volumes](#) on page 257

## Managing thin pool capacity

### Before you begin



- You can only reclaim thin pool capacity from bound thin volumes.
- Thin pool reclamation for individual thin volumes requires Enginuity 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.

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**
- Storage Groups (Enginuity 5876): **STORAGE > Storage Groups**
- Device Groups: **DATA PROTECTION > Device Groups**
- File Storage Groups: **SYSTEM > eNAS > File Dashboard > File Storage Groups**

### Procedure

1. Select the storage system.
2. Select **Storage > Volumes**.
3. Select the volume type by selecting a tab.
4. Do one of the following:
  - To start freeing unused capacity:

- - Select one or more volumes, click , and select **Start Allocate/Free/Reclaim** to open the **Start Allocate/Free/Reclaim** dialog box.
  - **Select Free Volumes.**
  - 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.
  - 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.
  - To stop freeing unused capacity:
    - - Select one or more volumes, click , and select **Stop Allocate/Free/Reclaim** to open the **Stop Allocate/Free/Reclaim** dialog box.
    - **Select Free Volumes.** In addition, on storage systems running Enginuity 5876, 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 5876.
    - 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.
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 920 and [Previewing jobs](#) on page 920.
  - 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*.

## Allocate/Free/Reclaim dialogs

Use the dialogs to perform the following operations:

- Start allocating thin pool capacity for thin volumes, as described in [Managing thin pool allocations](#) on page 244.
- Start freeing unused allocated thin pool capacity, as described in [Managing thin pool capacity](#) on page 253.
- Start reclaiming unwritten tracks from thin volumes, as described in [Managing space reclamation](#) on page 514.
- Stop allocating thin pool capacity for thin volumes, as described in [Managing thin pool allocations](#) on page 244.

- Stop freeing unused allocated thin pool capacity, as described in [Managing thin pool capacity](#) on page 253.
- Stop reclaiming unwritten tracks from thin volumes, as described in [Managing space reclamation](#) on page 514.

## Creating or Expanding or Modifying thin pools

### Advanced Options when creating thin pools

- Select the **Disk Group** containing the DATA volumes to use in the pool.
- 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 5876 or higher.
- 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 5876 or higher.
- To specify the percentage of the pool's capacity to enable, select **Enable Max Subscription (0-65534)** and type a percentage.
- To specify the percentage of the capacity of the thin pool that will be reserved for non-FAST activities, select **Enable 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.
- To enable the DATA volumes in the pool for use, select **Enable DATA Volume for Use**.
- 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.
- Click **APPLY**.

### Advanced Options when expanding thin pools

- Select **Start Write Balancing**.
- Click **APPLY**.

### Modifying thin pools

- Select **Volume Capacity**, in GB.
- 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 5876 or higher.
- 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 5876 or higher.

- To specify the percentage of the pool's capacity to enable, select **Enable Max Subscription (0-65534)** and type a percentage.
- To specify the percentage of the capacity of the thin pool that will be reserved for non-FAST activities, select **Enable 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.
- To enable the DATA volumes in the pool for use, select **Enable DATA Volume for Use**.
- 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.
- Click **OK**.

## Creating thin volumes

This procedure explains how to create thin volumes on storage systems running Enginuity version 5876. For instructions on creating thin volumes on storage systems running HYPERMAX OS 5977 or higher, refer to [Creating thin volumes](#) on page 185.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**, click on the **Virtual** tab and select **Create**.
3. Select **Configuration (TDEV or BCV + TDEV or Virtual Gatekeeper)**.
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. 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).
7. Click **Advanced Options** to continue setting the advanced options

Setting Advanced options:

- a. 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 196.

- b. To **Allocate Full Volume Capacity**, select the option.
  - c. 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.
  - d. 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.
  - e. 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.
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 920 and [Previewing jobs](#) on page 920.
  - Expand Add to Job List, and click **Run Now** to perform the operation now.
  - Click **Advanced Options** to continue setting the advanced options, as described next.

## Binding/Unbinding/Rebinding thin volumes

### Before you begin

This procedure applies to storage systems running Enginuity OS 5876.

- Only one bind, unbind, or rebind operation can be performed on the same volume in any one config session.
- 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*.

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.

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** and click on the **Virtual** tab.
3. Select the volume and do one of the following:
  - Click **FAST > Bind**
    - a. Select the thin pool with which to bind the volume.
    - b. Optional: Select **Allocate Full Volume Capacity** option.
    - c. To view additional information on the selected volumes, click Show selected volumes.
    - d. 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** option. Persistent allocations are unaffected by standard reclaim operations and any TimeFinder/Clone, TimeFinder/Snap, or SRDF copy operations.
    - e. Click **OK**.
  - Click **FAST > Unbind** and click **OK**
  - Click **FAST > Rebind**, specify the pool name, and click **OK**.

## Understanding Virtual LUN Migration

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

---

**Note**

Virtual LUN migration requires Enginuity 5876.

---

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

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

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

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

## Viewing VLUN migration sessions

### Procedure


1. Select the storage system.
2. Select **STORAGE > Vln Migration** to open the **Virtual LUN Migration** list view.

Use the this view to display and manage migration sessions.

The following properties display:


- **Name**—Migration session name.
- **Status**—Migration session status.
- **Invalid Tracks**—Number of invalid tracks for the volume pair.
- **Percentage**—Percentage of the session completed.

The following controls are available:

-  — [Viewing VLUN migration session details](#) on page 259
- **Terminate**—[Terminating a VLUN migration session](#) on page 260

## Viewing VLUN migration session details

### Procedure

1. Select the storage system.
2. Select **STORAGE > Vln Migration** to open the **Virtual LUN Migration** list view.
3. Select a session and click  to open its **Details** view.

Use this view to display details on a migration session. This view contains two panels: **Details** and **Source and Target Info**.

The following properties display in the **Details** panel:

- **Name**—Migration session name.
- **Status**—Migration session status.
- **Invalid Tracks**—Number of invalid tracks for the volume pair.
- **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.

The following properties display in the **Source and Target Info** panel:

- **Source**—Source volumes in the migration session.
- **Target**—Target volumes in the migration session.
- **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.

## Terminating a VLUN migration session

### Procedure

1. Select the storage system.
2. Select **STORAGE > Vln Migration** to open the **Virtual LUN Migration** list view.
3. Select the migration session and click **Terminate**.
4. Click **OK**.

## VLUN Migration dialog box

From this dialog box you can perform volume migrations for regular or thin volumes. Thin volumes migrate from a source pool to a target pool, and regular volumes migrate to configured (existing) volumes or unconfigured (new) volumes.

Some of the options in the dialog box display will differ depending on whether you are migrating regular or thin volumes.

For volume-specific migration procedures, refer to the following:

- [Migrating regular volumes](#) on page 261
- [Migrating thin volumes](#) on page 262
- [Migrating regular storage group volumes](#) on page 261
- [Migrating thin storage group volumes](#) on page 262

## Select VLUN Migration Session Target dialog box

Use this dialog box to select the target disk group (standard migration) or target thin pool (thin migration).

## Migrating regular storage group volumes



### Before you begin

- Virtual LUN migration requires Enginuity 5876.

This procedure explains how to migrate all the regular volumes in a storage group.

To migrate regular storage group volumes:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Groups** view.
3.  Select a storage group, click , and select **VLUN Migration**.
4. Type a **Migration Session Name**.  
Migration session names must be less than 32 characters long and are case sensitive.
5. Select the RAID **Protection type**.
6. Select **Target type**. Choose **Create new volumes** to migrate to unconfigured volumes or **Use existing volumes** to migrate to configured volumes.
7. Select whether to **Pin Volumes** so that they cannot be moved by any FAST automated process.
8. Click **OK** to create the migration session.

## Migrating regular volumes



### Before you begin

- Virtual LUN migration requires Enginuity 5876.

This procedure explains how to migrate individual regular volumes.

To migrate regular volumes:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes** to open the **Volumes** view.
3. Select the volume type by selecting a tab.
4.  Select one or more volumes, click , and select **VLUN Migration**.
5. Type a **Migration session name**.  
Migration session names must be less than 32 characters and are case sensitive.
6. Select the RAID **Protection type**.
7. Select the **Target type**. Choose **Create new volumes** to migrate to unconfigured volumes or **Use existing volumes** to migrate to configured volumes.
8. Select whether to **Pin Volumes** so that they cannot be moved by any FAST automated process.

9. Click **OK**.

## Migrating thin storage group volumes

### Before you begin


- Virtual LUN migration requires Enginuity 5876.

This procedure explains how to migrate all the thin volumes in a storage group.

To migrate thin storage group volumes:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups** to open the **Storage Groups** view.
- 3.

Select a storage group, click , and select **VLUN Migration**.

4. Type a **Migration Session Name**.

The session name must be less than 32 characters long and is case sensitive.

5. Select a **Target**.
6. From the **Migrate allocations from pool** menu, select a pool from which to migrate allocations.
7. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
8. Click **OK**.

## Migrating thin volumes

### Before you begin

- Virtual LUN migration requires Enginuity 5876.

This procedure explains how to migrate individual thin volumes.

To migrate selected thin volumes:

### Procedure

1. Select the storage system.
2. Select **STORAGE> Volumes** to open the Volumes view.
3. Select the volume type by selecting a tab.
- 4.

Select one or more thin volumes, click , and select **VLUN Migration**.

5. Type a **Migration Session Name**.

The session name must be less than 32 characters long and is case sensitive.

6. Select a **Target**.
7. From the **Migrate allocations from pool** menu, select a pool from which to migrate allocations.
8. Select whether to **Pin volumes** so that they cannot be moved by any FAST automated process.
9. Click **OK**.

# Understanding Federated Tiered Storage

Federated Tiered Storage (FTS) allows you to attach external storage to a storage system. Attaching external storage allows you to use physical disk space on existing storage systems while gaining access to features such as local replication, remote replication, storage tiering, data management, and data migration.

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*

## 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:

### Procedure

1. Select the storage system.
2. Select **STORAGE > External Storage**.

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, or multiple items in non-consecutive rows. 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.
- **Dir:Port**—Director: Port 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 LUNs 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 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**—[Virtualizing external LUNs](#) on page 264 (Only displays for Enginuity 5876)
- **Remove**—[Removing external LUNs](#) on page 266 (Only displays for HYPERMAX OS 5977 or higher)

## Virtualizing external LUNs

See [Virtualizing external LUNs](#) on page 265 for background information.

### Procedure

1. To virtualize external LUNs:
  1. Select the storage system.
  2. Select **STORAGE > External Storage**.
  3. (Optional) Click the **Not Virtualized** check box above the filtered LUNs 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.
  6. Select an import method from the **Import Method** drop-down menu. This determines the mode of operation for the eDisk.

#### **WARNING**

**Selecting Raw Space - External Provisioning deletes any data that is currently on the external volume.**

---



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 adds 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. Optional: Click **Advanced Options** to continue setting the advanced options, as described next.

Setting Advanced options:

- a. To override the auto meta member capacity configured on the storage 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.
- b. 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 is smaller than the other volumes in the meta.
- c. If you want to specify a DX director for the path to the eDisk, select a director from the **DX Director** drop-down menu.
- d. Click **OK**.

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Virtualizing external LUNs

When you attach external storage to a storage system, FAST.X virtualizes an external storage system's SCSI logical units as disks called eDisks. eDisks have two modes of operation:

### Encapsulation

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

### External Provisioning

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

The following restrictions apply to eDisks:

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

### Encapsulation

Encapsulation has two modes of operation:

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

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

**Encapsulation for virtual provisioning (VP encapsulation)**

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

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

**External provisioning**

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

---

**Note**

If you use external provisioning, any data that is 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 does not always match the raw capacity of the eDisk. If the capacity does not match, Enginuity sets a custom geometry on the encapsulated volume. For created meta volumes, Enginuity defines the geometry on the meta head, and only the last member can have a capacity that spans beyond the raw capacity of the eDisk.

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

## Removing external LUNs

**Before you begin**

- This feature requires HYPERMAX OS 5977 or higher.
- LUNs must be virtualized.

This procedure explains how to remove external LUNs from storage groups protected with ProtectPoint. Encapsulated LUNs whose volumes are in a storage group cannot be removed.

**Procedure**

1. Select the storage system.
2. Select **Storage > External Storage**.  
Opens 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). Select either a single item, multiple items in consecutive rows, or multiple items in non-consecutive rows. 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**.  
Opens 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:
  - **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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List** , and click **Run Now**.

## 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 provisioning storage groups.

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

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

A provisioning template contains configuration information and a performance reservation.

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

## Creating storage templates

### Before you begin


- Storage templates require HYPERMAX OS 5977 or greater.
- This feature is only available to a user with Admin or StorageAdmin permission.

Using the configuration and performance characteristics of an existing storage group as a starting point, you can create templates that will pre-populate fields in the provisioning wizard and create a more realistic performance reservation in your future provisioning requests.

To create a storage template:

### Procedure

1. Select the storage system.
2. Select **STORAGE > Templates** to open the **Provisioning Templates** list view.
3. For second and subsequent templates, click **Provision**. Go to step 7.
4. To create the first template, click **Select a Storage Group** (this is part of the text under **Get Started!**). and select a storage group that has a service level (SL) assigned (FBA only).

5. Click  and click the **Compliance** icon.
6. Click **Save as a Template**.
7. Review the default values and update as appropriate.

Configuration information includes the Service Level, Workload Type, Number and Size of Volumes to be saved as part of the template. By default the information will be populated based on the selected storage group, which can be modified as required before saving.

**Service Level:** The drop-down will be populated with the all the available service levels on the selected storage system and Array Default. By default, the service level of the selected storage group will be selected.

**Workload Type:** The drop-down will be populated with the workload types available to the selected SL (including None). By default the workload type of the selected storage group will be selected.

**Volumes:** By default the number of volumes will be the number of volumes in the selected storage group. This field can be left empty. If the Scale Limits switch is on (default state), a change made to volume size will scale the IOPS and MBPS chart and an appropriate host IO limit will be calculated as a recommended value.

**Volume Size:** By default the size of the volumes will be the size of the volumes in the selected storage group. If there are multiple volume sizes, the size of the first volume size encountered will be used. This field can be left empty. Volume capacity units available will be GB and TB. If the Scale Limits switch is on (default state), a change made to volume size will scale the IOPS and MBPS chart and an appropriate host IO limit will be calculated as a recommended value. Read Interaction Between Charts/Data for more information.

**Expected RT:** The expected average response time for the selected service level.

The Host I/O Limit section will pull current host I/O limit information from the source storage group (standalone or child limit only, parent limit is ignored). If no host I/O limit is set, a host I/O limit in IOPS is recommended and the value is pre-populated.

Host I/O Limit combobox:

Options: IOPS, MBPS, Both, and None.

Initial Values:

- If the source Storage Group has an IOPS limit set, **IOPS** is selected.
- If the source Storage Group has an MBPS limit set, **MBPS** will be selected.
- If the source Storage Group has an IOPS and an MBPS limit set, **Both** is selected.
- If the source Storage Group has no limit set, **IOPS** is selected.

Host I/O Limit input field(s) and associated recommendation label (s):

- IOPS Selected in combobox:
  - Text input initial value:
    - If source storage group has an IOPS limit set, initial value is that limit.
    - If source Storage Group has no IOPS limit set, initial value is the recommended limit.

- Restrictions:
  - Value must fall between 100-2,000,000.
  - Value must be a multiple of 100.
- MBPS selected in combobox.
  - Text input initial value:
    - If source storage group has an MBPS limit set, initial value is that limit.
    - If source Storage Group has no MBPS limit set, initial value is the recommended limit.
    - Restrictions:
      - Value must fall between 1-100,000.
- Both selected in combobox - the IOPS and MBPS information is displayed.
- None selected in combobox - no text field or recommendation is displayed.

Scale limits switch: The scale limits switch is enabled by default.

If the switch is 'on', provisioning requests using this template with scale the host IO Limit recommendation(s) if the template's default capacity is overridden.

If the switch is 'off', provisioning requests will use the exact Host IO Limit value(s) that were saved with the template.

Dynamic Distribution:

- Options: Never, OnFailure, and Always.
- Initial value:
  - If source Storage Group has a Dynamic Distribution value, initial value is that value.
  - If source Storage Group has no Dynamic Distribution set, initial value is Never.

Performance Reservation

This is a 2 week expiring performance reservation or plan that will be used for comparison on the storage group details page and suitability checks.

IOPS and MBPS: Similar to IOPS and MBPS chart on the storage group details page with just the actual values and no plan (more details Workload Compliance Details spec).

A Host I/O Limit line will be seen on the graph if a corresponding value has been set in the Host IO Limit section.

Workload Skew: Similar to the workload skew chart on the storage group details page with just the actual values and no plan.

I/O Mixture: Similar to the workload mixture chart on the storage group details page with just the actual values and no plan.

Interaction Between Charts and Data

When Capacity is Modified (both Volumes Count and Size are Populated)

Performance Reservation Section:

IOPS/MBPS: The IOPS and MBPS values are scaled to the new capacity to preserve the IO density. For example, if the source storage group's total capacity was 10 x 50 GB volumes = 500GB, and the Volumes field was changed from 10 to 15, the total capacity would be 750GB. The 42 IOPS values and the

42 MBPS values in the charts would be multiplied by 1.5 to reflect the 50% increase in capacity.

I/O Mixture: No change.

Skew: No change.

If the Scale limits switch is 'on', the recommended limit will be recalculated according to the new capacity. So it would be two times the maximum 42 bucket value of IOPS and/or MBPS as calculated for the Performance Reservation section. The textbox will be auto-populated with the recommended limit. This new value will also be drawn in the appropriate IOPS and/or MBPS chart in the Performance Reservation section.

If the Scale limits switch is 'off', the recommended limit will be recalculated the same way. The value in the textbox will NOT be overwritten in this case.

When Capacity is Modified (Volumes Count and/or Size are Empty)

Performance Reservation Section:

IOPS/MBPS: Total capacity it required to calculate IOPS/MBPS. If we are missing count, size, or both, the IOPS and MBPS values will be calculated according to the current total capacity of the source storage group. So if (for example) the storage group is 5 x 100GB devices, and Volume Count is nulled out, and device size is change to 75GB, IOPS and MBPS will be calculated assuming 500GB. The assumed capacity will be displayed in the upper right hand corner of the charts.

I/O Mixture: No change.

Skew: No change.

Host IO Limit:

Total capacity it required to calculate Host IO Limit recommendation. If we are missing count, size, or both, the recommendation will be calculated according to the current total capacity of the source storage group. An information icon will be shown next to the recommendation. Hovering will give more information.

When Host IO Limit Combobox and Textbox Values are Modified

The specified value will be drawn in the appropriate charts. If the value is updated, the chart is updated.

If the combobox value is Both, you should have a red line corresponding to the specified value of IOPS on the IOPS chart and MBPS on the MBPS chart. If None is selected, no host IO limit line should show up on either IOPS or MBPS chart.

If IOPS is selected, there should be a red line on IOPS and none on MBPS. If MBPS is selected, there should be a red line on MBPS and none on IOPS.

#### 8. Click **SAVE**

If there has not been at least one week of data collected for the selected storage group, a dialog is displayed (see [Dialog displayed when there is less than one week's data collected](#) on page 96).

## Viewing storage templates

### Before you begin

- Storage templates require HYPERMAX OS 5977 or greater.

The **Provisioning Template** list view allows you to view and manage provisioning templates.

## Procedure

1. Select the storage system.
2. Select **STORAGE > Templates** to open the **Provisioning Template** list view.
3. Select a template card.

The following properties are displayed - template service level, workload type, response time, capacity information (number of volumes, size and headroom), as well as workload characteristics (I/O density, I/O size, Writes and Skew.)

4. Hover near the workload writes % to view a popup chart of the I/O mixture that the workload is running. Hovering over the sections of the pie chart reveals the percentages associated with each I/O type. To dismiss the popup charts simply click anywhere off of the charts.
5. Hover near the workload skew % to view a popup chart of the actual workload skew.

The actual workload skew is a load percentage over the percentage of capacity used in the workload. Hovering over the line on the chart will display the percentages for actual capacity and load score. To dismiss the popup charts simply click anywhere off of the charts.

6. Click on the icon on the top-right hand corner of the template card to view the back of the template card.

The back of the template card displays the name of the template at the very top along with two charts underneath it. The top chart displays the set workload host IO limit in IOPS along with the actual workload IO statistics. The bottom chart displays the set workload host IO limit in MBPS along with the actual workload MBPS IO statistics.

7. Click on the icon on the top-right hand corner of the template card to view front side of the card again.

The following controls are available:

- **Provision**— [Creating storage templates](#) on page 267
- **Modify**— [Modifying storage templates](#) on page 271
-  — [Deleting storage templates](#) on page 272

## Modifying storage templates

### Before you begin

- Storage templates require HYPERMAX OS 5977 or greater.
- The user must have Administrator or StorageAdmin permission.

### Procedure

1. To modify a storage template:
  1. Select the storage system.
  2. Select **STORAGE > Templates** to open the **Provisioning Templates** list view
  3. Select the template and click **Modify** to open the **Modify Template** wizard.
  4. Modify the template as you step through the wizard.

All of the fields are exactly like the **Save as a Template** dialog. The only difference is how the scale limits work when volume size or volume count field is left empty.

If the selected template has both volume size and count when it was created and the user removes one or either of them during the modification operation with scale limits switch on, the original capacity of the template will be used for the display purposes and it will be shown in the tooltip next to the IOPS/MBPS field and top right corner of the chart.

If the selected template did not have either volume size and count when it was created and the user leaves one or either of the fields empty modification operation with scale limits switch on, the capacity of 200 GB will be used for the display purposes and it will be shown in the tooltip next to the IOPS/MBPS field and top right corner of the chart.

5. Click **Finish**.


## Deleting storage templates

### Before you begin

- Storage templates require HYPERMAX OS 5977 or greater.
- This feature is only available for a user with Administrator or StorageAdmin permission.

To delete a storage template:

### Procedure

1. Select **STORAGE > Storage Templates** to open the **Storage Template** list view.
2. Select the template and click  .
3. Click **OK**.

## Understanding FAST.X

FAST.X allows the seamless integration of storage systems running HYPERMAX OS 5977 or higher 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*

## Viewing external disks

### Before you begin



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.



## Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Resource Pools**.
3. Select the SRP and click  to view its details.
4. Click the number next to **Disk Groups**.
5. Select an external disk group and click  to view its details.
6. Click the number next to **Number of disks**.

The following properties are displayed:

### Name

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.

### 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](#) on page 274
- **Remove** — [Removing external disks or External LUNs](#) on page 274
- **Start Draining** — [Start draining external disks](#) on page 275
- **Stop Draining** — [Stop draining external disks](#) on page 276
- **Activate** — [Activating external disks](#) on page 276

## Adding external disks

### Before you begin

This action can be performed only for HYPERMAX OS 5977 or higher.

You can add an external disk to the external disk group of a storage resource pool (SRP).

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.

### Procedure

1. To add an external disk:
  1. Select the storage system.
  2. Select **Storage > Storage Resource Pools**.
  3. Select the SRP.
  4. Click **Add eDisks**.

The **Add eDisks** dialog box shows the available external LUN WWNs from multiple external arrays.

5. Select the external disk to be added.
6. If you want to preserve the existing data on the external LUN, select **Incorporate eDisk data**. If you leave the **Incorporate eDisk data** cleared, the existing data on the external LUN is cleared.
7. (Optional) In the **Add Storage Group** list, select a storage group to add.

You can filter the list by searching for a storage group by name.

This option is available only on storage systems running HYPERMAX OS 5977 Q1 2016.

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Removing external disks or External LUNs

### Before you begin

See [Removing external LUNs](#) on page 266 for information on removing external LUNs.



See below for information on removing external disks.


This action can be performed only for HYPERMAX OS 5977 or higher.


You can remove an external disk from a storage resource pool (SRP) if it is in a Drained state.

### Procedure

1. To remove an external disk:

1. Select the storage system.
2. Select **STORAGE > Storage Resource Pools**.
3. Select the SRP and click  to view its details.
4. Click the number next to **Disk Groups**.
5. Select an external disk group and click  to view its details.
6.
 



 Select the external disk that you want to remove, click , and click **Remove eDisks**.  
 The **Remove External LUNs** dialog appears 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Working with external disks

You can perform the following operations:

- Start draining an external disk. For more information, refer to [Start draining external disks](#) on page 275.
- Stop draining an external disk. For more information, refer to [Stop draining external disks](#) on page 276.
- Activate an external disk. For more information, refer to [Activating external disks](#) on page 276.



## 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.

### Procedure

1. To start a drain operation on an external disk:
  1. Select the storage system.
  2. Select **STORAGE > Storage Resource Pools**.
  3. Select the SRP and click  to view its details.
  4. Click the number next to **Disk Groups**.
  5. Select an external disk group and click  to view its details.

6.

Select the external disk that you want to drain, click , and click **Start Draining**.

7. 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 920 and [Previewing jobs](#) on page 920.
- 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.


### Procedure

1. To stop a draining operation on an external disk:

1. Select the storage system.

2. Select **STORAGE > Storage Resource Pools**.3. Select the SRP and click  to view its details.4. Click the number next to **Disk Groups**.5. Select an external disk group and click  to view its details.

6.

Select the external disk that you want to stop draining, click , and click **Stop Draining**.

7. 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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Activating external disks

### Before you begin


This action can be performed only for HYPERMAX OS 5977 or higher.



You can activate an external disk if it is in a draining, drained, or disabled state.

### Procedure

1. To activate an external disk:

1. Select the storage system.

2. Select **STORAGE > Storage Resource Pools**.3. Select the SRP and click  to view its details.4. Click the number next to **Disk Groups**.5. Select an external disk group and click  to view its details.

6.  Select the external disk that you want to activate, click , and click **Activate**.
7. 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing reservations

### Procedure

1. Select the storage system.
2. In the **Dashboard**, click the **System Health** tab.
3. In the **Actions** panel, click **View Reservations**.

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. The default value is **Never**.
- **User Comment**—User-supplied comments.

The following control is available:

**Release**—[Releasing reservations](#) on page 278

## Viewing reservation details

### Procedure

1. Select the storage system.
2. In the dashboard, click the **System Health** tab.
3. In the **Action** panel, click **View Reservations**.
4. Select the reservation and click .

The **Properties** panel displays the following:

- **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.

There are 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** will open a view listing the volumes held in the reservation.

## Releasing reservations

### Procedure

1. Select the storage system.
2. In the dashboard, click the **System Health** tab.
3. In the **Action** panel, click **View Reservations**.
4. Select one or more reservations and click **Release**.
5. Click **OK**.

## Managing vVol

### Before you begin

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

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

To access the VVol Dashboard:

### Procedure

1. Select the storage system.
2. Select **Storage > VVol Dashboard**.

VVol Dashboard:

The VVol Dashboard is organized into the following panels:

#### Summary panel

Displays the following VVol summary information:

- **Storage Containers** — The number of storage containers on the selected storage system.  
Click **Storage Containers** to display the Storage Containers list view. For more information about viewing storage containers, refer to [Viewing storage containers](#) on page 279.
- **Protocol Endpoints** — The number of protocol endpoints on the selected storage system.  
Click **Protocol Endpoints** to display the Protocol Endpoints list view. For more information about protocol endpoints, refer to [Viewing protocol endpoints](#) on page 285.
- **PE Masking Views** — The number of masking views that contain protocol endpoints.  
Click **PE Masking Views** to display the PE Masking Views list view. For more information about PE masking views, refer to [Viewing masking views](#) on page 308.

To view additional information on a particular item, click on it to open the corresponding list view.

#### Actions panel

Displays links to the following common tasks:





- **CREATE STORAGE CONTAINER** — [Creating storage containers](#) on page 281
- **PROVISION PROTOCOL ENDPOINT TO HOST** — [Provisioning protocol endpoints to hosts](#) on page 286
- **STORAGE CONTAINER ALERTS** — [Viewing alerts](#) on page 52

#### Symmetrix Consumed Capacity - Subscribed panel

Displays a bar graph representing how much subscribed space all storage containers consume on the storage system.

#### VASA Provider Status panel

Displays one of the following icons representing the status of the VASA provider:

-  — The VASA provider is online.
-  — The VASA provider is offline.
-  — A connection to the VASA provider has not been configured.
-  — There was an error connecting to the VASA provider.

To refresh the status of the VASA provider, click .

To create a connection to the VASA provider, click **Create Connection**. To edit an existing connection, click **Edit Connection**. For more information about configuring a connection to the VASA provider, see [Configuring the VASA provider connection](#) on page 287.

#### Storage Resources panel

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.
- **Compression** — If compression is enabled on this storage resource a tick will appear. If it's disabled a horizontal dash will appear.

Click **VIEW ALL STORAGE RESOURCE** to view the **Storage Resources** list view.

## Viewing storage containers

To view the storage container list:


**Procedure**

1. Select the storage system.
2. Select **Storage > VVol 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:

-  — [Viewing storage container details](#) on page 280
- **Create**—[Creating storage containers](#) on page 281
- **Modify**—[Modifying storage containers](#) on page 282
- **Delete**—[Deleting storage containers](#) on page 282

## Viewing storage container details

To view storage container details:

**Procedure**

1. Select the storage system.
2. Select **Storage > VVol Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container and click .

The following properties display:

- **Name** — The name of the storage container.
- **Description** — The description of the storage container. This field is editable.
- **Subscribed Limit (GB)** — The total combined limit of all storage resources within the storage container.
- **Subscribed Used (GB)** — The current subscribed usage on the storage container of all of the storage resources within the storage container.
- **Subscribed 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 Resources** — The total number of storage resources within the storage container.



## Creating storage containers

This procedure allows you to create a storage container. To add a storage resource to a storage container, refer to [Adding storage resources to storage containers](#) on page 284.


To create a storage container:

### Procedure

1. Select the storage system.
2. Select **Storage > VVol Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Click **Create** . The **Create Storage Container** wizard displays.
5. 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  and complete the following steps (same steps when modifying an existing resource):

- a. In the **Name** field, type a name for the storage resource, or accept the default name.
  - b. From the **SRP** menu, select the SRP to apply to the storage resource.
  - c. From the **Service Level** menu, select the service level to apply to the storage resource.  
For all-flash storage systems, the only service level available is Diamond and it is selected by default.
  - d. From the **Workload** menu, select the workload to apply to the storage resource.
  - e. In the **Limit (GB)** field, type the imposed subscribed limit on the storage resource. 0.1 GB is the minimum value allowed.
8. Compression is enabled by default on All Flash systems running the HYPERMAX OS 5977 Q3 2016 Service Release or higher. To disable the feature on this storage container, uncheck the **Compression** check box. For more information, refer to [Understanding compression](#).
  9. Click **NEXT**.
  10. On the **Summary** 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 920 and [Previewing jobs](#) on page 920
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Modifying storage containers

To modify a storage container:

### Procedure

1. Select the storage system.
2. Select **Storage > VVol Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container and click **Modify**.
5. Modify the description.
6. Click **OK**.

## Deleting storage containers


### Before you begin

- The storage system must be running HYPERMAX OS 5977 or higher.
- You cannot delete containers with used capacity.

To delete a storage container:

### Procedure

1. Select the storage system.
2. Select **Storage > VVol Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
- 4.


Select the storage container you want to delete, click , and click **Delete**.

5. Click **OK**.

## Viewing storage resources

To view the storage resource list:

### Procedure


1. Select the storage system.
2. Select **Storage > VVol Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container and click .
5. Click the number next to **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.
- **Service Level** — The name of the service level.
- **Workload** — The name of the workload.
- **Subscribed Used (%)** — The current percent of subscribed tracks within the storage resource in relation to the limit imposed on the storage resource.

- **Subscribed Limit (GB)** — The subscribed capacity limit within the storage resource.



The following controls are available:

-  — [Viewing storage resource details](#) on page 283
- **Add** — [Adding storage resources to storage containers](#) on page 284
- **Modify** — [Modifying storage resources](#) on page 285
- **Remove** — [Removing storage resources from storage containers](#) on page 285

## Viewing storage resource details

To view storage resource details:

### Procedure

1. Select the storage system.
2. Select **STORAGE > VVol Dashboard**.
3. Click **Storage Containers**.
4. Select the storage container and click .
5. Click the number next to **Storage Resources** to display the Storage Resources list view.
6. Select the storage resource and click .


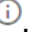
The following properties display:

- **Name** — The name of the storage resource.
- **Storage Container** — The name of the associated storage container.
- **SRP** — The name of the associated SRP.
- **Service Level** — The name of the associated service level.
- **Workload** — The name of the associated workload.
- **Compression** — Indication if compression is enabled or disabled.
- **Compression Ratio** — The current compression ratio on this storage resource.
- **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.

## Viewing storage resource related SRPs

To view the related SRPs of a storage resource:

**Procedure**


1. Select the storage system.
2. Select **Storage > VVol Dashboard**.
3. Click **Storage Containers** to display the Storage Containers list view.
4. Select the storage container and click .
5. Click the number next to **Number of Storage Resources** to display the Storage Resources list view.
6. Select the storage resource and click .
7. Click the entry next to SRP to display the Storage Resource Pools list view.

For more information about the Storage Resource Pools list view, refer to [Viewing Storage Resource Pools](#) on page 154.


**Adding storage resources to storage containers**


To add a storage resource to a storage container:



**Procedure**

1. Select the storage system.
2. Select **STORAGE > VVol Dashboard**.
3. Click **Storage Containers** to display the Storage Containers list view.
4. Select the storage container and click .
5. Click the number next to **Number of Storage Resources**.
6. Click **Add**.

The **Add Storage Resource To Storage Container** dialog box displays. The details of any existing storage resource are populated automatically.

7. To add an additional resource, click  and specify the following details:
  - **Name** — The name of the storage resource.
  - **SRP** — The name of the SRP.
  - **Service Level** — The name of the service level.
  - **Workload** — The name of the workload.

For more information about current workload, click .


  - **Limit (GB)** — The subscribed capacity limit imposed.
  - **Compression** — The **Compression** check box will be checked if you enabled compression when creating the storage group. Uncheck to disable compression on this particular storage resource. For more information, refer to [Understanding compression](#).
8. (Optional) If required, edit the details of the new storage resource, click  to remove it completely, or click  to add another new storage resource.
9. After you have added all of the required storage resources, 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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Modifying storage resources

To modify a storage container:



### Procedure

1. Select the storage system.
2. Select **STORAGE > VVol Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container and click .
5. Click the number next to **Number of Storage Resources**.
6. Select the storage resource and click **Modify**.
7. Modify the subscribed limit.
8. Click **OK**.

## Removing storage resources from storage containers

To remove a storage resource from a storage container:

### Procedure

1. Select the storage system.
2. Select **STORAGE > VVol Dashboard**.
3. Click **Storage Containers** to display the **Storage Containers** list view.
4. Select the storage container and click .
5. Click the number next to **Number of Storage Resources**.
6. Select the storage resource you want to remove, click , and click **Remove**.
7. Click **OK**.



## Viewing protocol endpoints

To view the protocol endpoints list:

### Procedure



1. Select the storage system.
2. Select **Storage > VVol 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:

-  — [Viewing protocol endpoint details](#) on page 286
-  — [Deleting protocol endpoints](#) on page 287
- **Set Volume Identifier** — [Setting volume identifiers](#) on page 196

## Viewing protocol endpoint details

To view protocol endpoint details:

### Procedure

1. Select the storage system.
2. Select **Storage > VVol Dashboard**.
3. Click **Protocol Endpoints**.
4. Select the protocol endpoint and click .

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.

## Provisioning protocol endpoints to hosts

To provision a protocol endpoint to a host:

### Procedure

1. Select the storage system.
2. Select **Storage > VVol Dashboard**.
3. In the **Actions** panel, click **Provision Protocol Endpoint to Host**.
4. 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](#) on page 292.
  - To create a host group, click **Create Host Group**. The **Create Host Group** dialog displays. For more information, refer to [Creating host groups](#) on page 302.

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 316.
  - To use an existing port group, select **Existing**, and select a port group from the **Port Group** list.
7. Click **NEXT**.
8. On the **Summary** page, review the details and do one of the following:
  - Optional: Modify the auto-generated **Masking View** name.
  - Optional: Modify the auto-generated **Storage Group** name.
  - 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.


## Deleting protocol endpoints

### Before you begin

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

To delete a protocol endpoint:

### Procedure

1. Select the storage system.
2. Select **Storage > VVol Dashboard**.
3. Click **Protocol Endpoints** to display the Protocol Endpoints list view.
4. Select the protocol endpoint you want to delete and click .
5. Click **OK**.

## Configuring the VASA provider connection

To configure the VASA provider connection:

### Procedure

1. Select the storage system.
2. Select **Storage > VVols Dashboard**.
3. In the **VASA Provider Status** panel, do one of the following:
  - To create a new connection, click **Create Connection**.
  - To edit an existing connection, click **Edit Connection**.
4. Specify the IP address of the VASA provider.
5. Click **OK**.

## Understanding compression

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

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

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

The following are the prerequisites for using compression:

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

### Reporting

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

---

### Note

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

---

## Viewing the SRP efficiency details

### Before you begin

Users need to have at least Monitor rights.

This procedure explains how one way to view the overall efficiency details of an SRP. The **Overall Efficiency Ratio** field can also be viewed from the **Storage Resource Pools Details** view.

### Procedure

1. Select the storage system.
2. Select **CAPACITY** to open the **CAPACITY** dashboard.

The following fields are displayed in the **Efficiency** panel:

- Overall Efficiency Ratio- The ratio of the sum of all TDEVs and Snapshot sizes and the Physical Used Storage (calculated based on the compressed pool track size).
- Virtual provisioning Savings - The ratio of the sum of all TDEVs and Snapshot sizes and the sum of all TDEVs allocated plus RDP allocated space.
- Snapshot Savings - The ratio of the RDP Logical Backend Storage (calculated based on the 128K track size) and the RDP Physical Used Storage of the RDP space (calculated based on the compressed pool track size).

## Viewing compressibility reports

This procedure shows how to view maximum data compressibility of storage groups on an All Flash storage system. Compression must be enabled on the storage system.



Before you begin:

- This feature requires HYPERMAX OS 5977.1125.1125 running on an All Flash storage system
- The account you use on Unisphere must have Monitor privilege at least.

#### Procedure

1. Select the storage system.
2. Select **CAPACITY** to open the **CAPACITY** dashboard.
3. Select a SRP instance from the drop down menu and in the **Actions** panel, click **COMPRESSIBILITY**.

The report lists the following details for each storage group:


- **Storage Group**—The name of the storage group.
- **# of Volumes**—The number of volumes in the group.
- **Allocated (GB)**—The amount of space allocated to the storage group.
- **Used (GB)**—The amount of allocated space that the group is using.
- **Target Ratio**—The expected compression ratio based on the last 24 hours of samples. If all storage groups are compressed, the compressibility report will be empty except for an entry named **NOT\_IN\_SG** (assuming that not all of the configured volumes are in storage groups).

## Viewing a storage group's compression ratio

#### Before you begin

Users need to have at least Monitor rights to view the compression ratio.

#### Procedure

1. Select a storage system.
2. Select **STORAGE > Storage Groups**
3. Select a storage group and click .

The **Compression**, **Compression Ratio** and **VP Saved** fields for the selected storage group are displayed.

If compression is enabled on the storage group a tick will appear in the **Compression** field. If compression is disabled a horizontal dash will be shown.


## Viewing a volume's compression details

#### Before you begin

Users need to have at least Monitor rights to view the compression ratio.

This procedure explains how to view a storage group volume's compression ratio.

#### Procedure

1. Select a storage system.
2. Select **STORAGE > Storage Groups**.
3. Click  and click the number next to **Volumes**.

4. Select a volume and click .

The **Compression Ratio** field for the selected volume is displayed. If compression ratio is not applicable on the volume the field will read "N/A."

5. Alternatively, select a storage system and then select **Storage > Volumes**.

6. Select a volume and click .

The **Compression Ratio** field for the selected volume is displayed.

## Viewing compression status using the VVol Dashboard

### Before you begin



Users need to have at least Monitor rights.

This procedure explains how to view the compression status and compression ratio of storage resources using the **VVol Dashboard**.

### Procedure

1. Select a storage system.
2. Select **Storage > VVol Dashboard**.

The compression state column for each storage resource is displayed in the **Symmetrix Consumed Capacity - Subscribed** panel. If compression is enabled for that resource a tick will appear in the column. If compression is disabled a horizontal dash will be shown.

3. To view the compression ratio on a storage resource, click on **Storage Containers**.
4. Select a storage container and click .
5. Click the number next to **Storage Resources**.
6. Select a storage resource and click .

The **Compression Ratio** field is displayed.

## Viewing the compression efficiency dashboard

This procedure explains how to view the compression efficiency of a storage system running HYPERMAX OS 5977.

### Procedure

1. Select a storage system.
2. Select **PERFORMANCE > Dashboards**.
3. Choose **Array** as the category.
4. Click on the **Array Efficiency** tab.

# CHAPTER 5

## Host Management

• Understanding Host Management.....	292
• Creating hosts.....	292
• Creating host groups.....	302
• Creating masking views.....	307
• Setting initiator port flags.....	311
• Setting initiator attributes.....	312
• Renaming initiator aliases.....	312
• Replacing initiators.....	313
• Removing masking entries.....	313
• Viewing initiators.....	313
• Viewing initiator details.....	314
• Viewing volumes associated with host initiator.....	315
• Viewing details of a volume associated with initiator.....	316
• Creating port groups.....	316
• Managing storage for Mainframe.....	322
• Mapping CKD volumes.....	340
• Creating PowerPath hosts.....	343
• Viewing PowerPath hosts.....	344
• Viewing PowerPath hosts details.....	344
• Viewing PowerPath Host Virtual Machines.....	345
• Viewing host cache adapters.....	346

# Understanding Host Management

Host Management covers the following areas:

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

## Creating hosts

### Before you begin

- To perform this operation, you must be a StorageAdmin.
- The storage system must be running Enginuity version 5876, or HYPERMAX OS 5977 or higher.
- The maximum number of initiators allowed in a host depends on the storage operating environment:
  - For Enginuity 5876, the maximum allowed is 32.
  - For HYPERMAX OS 5977 or higher, the maximum allowed is 64.

To create hosts:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3. Click **Create > Create Host**.  
The Create Host dialog displays.
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 hyphens ( - ) 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.

Select a host, click  and then click **Set Flags** to open the **Set Host/Host Group Flags** dialog.

7. Optional: To set the host port attributes:

- a. Click **Set Host Flags**.
- 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**.

8. Do either of the following:

- Click **Run Now** to start the task now.
- 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 920 and [Previewing jobs](#) on page 920.

## Adding initiators to hosts

### Before you begin

To perform this operation, you must be a StorageAdmin.

The storage system must be running Enginuity version 5876 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:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3. Select the host and click **Modify** to open the **Modify Host** dialog.
- 4.

Select an initiator from the **Available Initiators** list and click .

5.

To add a user defined initiator to the host, click , fill in the name and click **OK**.


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 **Run Now** or **Add To Job List**.

## Adding initiator to host

To add an initiator to a host:

### Procedure

1. Select **HOSTS > Hosts**
2. Click **Create** and then click **Create Host**.
3. Click the  button to the right of **Initiators in Host**.
4. Type the Initiator name.
5. 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 5876 or higher.

To remove initiators from hosts:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3. Select the host and click **Modify** to open the **Modify Host** dialog.
- 4.

Select an initiator from the **Available Initiators** list and click  .

5. Select the initiator and click **Replace Initiator**.
6. Click **Run Now** or **Add To Job List**.

## Modifying hosts

### Before you begin

To perform this operation, you must be a StorageAdmin.

The storage system must be running Enginuity version 5876 or higher.

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts** to open the Hosts list view.
3. Do one of the following:
  - Modifying hosts:
    - Select the host and click **Modify** to open the **Modify Host** dialog box.
    - 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.

- Adding initiators:
  - 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.


- 

Select an initiator from the **Available Initiators** list and click .

- Repeat these steps for each additional initiator.
- To add a user defined initiator to the host, click , fill in the name and click **OK**.

- Removing initiators:

- 

In the **Initiators in Host** list, select the initiator and click .

- Repeat these steps for each additional initiator.

#### 4. Do either of the following:

- Click **Run Now** to start the task now.
- 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 920 and [Previewing jobs](#) on page 920.

## Renaming hosts/host groups

### Before you begin

To perform this operation, you must be a StorageAdmin.

The storage system must be running Enginuity version 5876 or higher.

To rename host/host groups:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3. Select the host/host group and click **Modify**.
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 or host group port flags

To set host or host group port flags:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3.
 



Select a host, click  and then click **Set Flags** to open the **Set Host/Host Group Flags** dialog.
4. Optional: Select a host/host group whose flag settings you want to copy from the **Copy Flags from Other Host/Host Group** drop-down menu.
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 5876 or higher.

To delete hosts/host groups:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3.
 



Select the host/host group from the list and click .
4. Click **Delete**, then click **OK** to confirm.

## Viewing hosts/host groups

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts** to open the Hosts list view.
3. 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.

Click  to view the host/host group details.

The following controls are available:

**Create Host**—[Creating hosts](#) on page 292

**Create Host Group**—[Creating host groups](#) on page 302

**Provision Storage to Host**—[Using the Provision Storage wizard](#) on page 100 or [Using the Provision Storage wizard](#) on page 108


**Modify**—[Modifying hosts](#) on page 294 or [Modifying host groups](#) on page 304

**Set Flags**—[Setting host or host group port flags](#) on page 296

**Delete**—[Deleting hosts/host groups](#) on page 296

## Viewing host/host group details

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3. Select the host/host group.
4. Click  to view the host/host group details.

---

### 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.

**Hosts**—Number of hosts in the group. This field only displays for host groups.

**Masking Views**—Number of masking views with which the host/host group is associated.

**Initiators**—Number of initiators in the host/host group. For host groups, the value includes initiators in any child host groups.

**Host Groups**—Number of host groups in which this host is a member. This field only displays for individual hosts. .

**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.

**PowerPath Hosts**—Number of PowerPath hosts.

## Viewing host initiators

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
- 3.

Select the host and click  to open the host details view, then click on the link in the **Initiators** field to open the initiators list view.

The following properties display:

**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.

**Hosts**— Number of hosts the initiator is associated with.

**Masking Views**—Number of associated masking views.

The following controls are available:

**Set Attributes**—[Setting initiator attributes](#) on page 312

**Set Host Flags**—[Setting initiator port flags](#) on page 311

**Rename Alias**—[Renaming initiator aliases](#) on page 312

**Replace Initiator**—[Replacing initiators](#) on page 313

## Remove Masking Entry—[Removing masking entries](#) on page 313

### 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 5876.
Volume Set Addressing**	<p>Enables the volume set addressing mode.</p> <p>When using volume set addressing, you must specify a 4-digit address in the following range:</p> <p>(0)000-(0)007, (0)010-(0)017,... to a maximum of (0)FF0-(0)FF7</p> <p>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.</p>
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).
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)*	<p>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.</p> <p>To enable the SCSI Support1 attribute, you must also enable the SPC2 Protocol Version attribute.</p>
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 systems running Enginuity 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:

Clear the **Override** option and click **OK** to close the dialog.

Open the dialog again, select **Override**, and then the desired state (Enable/Disable).

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 I/O limits dialog box

Use this dialog box to set the host I/O limits for the storage group you are provisioning:

### Procedure

1. Type values for one or both of the following:
  - **MB/Sec**—Maximum bandwidth (in MB per second). Valid values range from 1 MB/sec to 100,000 MB/sec.
  - **IO/Sec**—Maximum IOPs (in I/Os per second). Valid values range from 100 IO/Sec to 2,000,000 IO/sec, in 100 increments.
2. 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.
  - **Failure**—Enables port failure capability. When enabled, the fraction of configured host I/O limits available to a configured port will adjust based on the number of ports currently online.
3. Click **OK**.

---

### Note

For more information on host I/O limits, refer to [Setting host I/O limits](#) on page 132.

---

## Host Group filtering rules

The host and host group list follows these guidelines for display:

<p>Initiators with the same name, but seen from different storage system login history tables will be filtered to only show once. New host groups can be set on both storage systems.</p>
---

<p>Initiators logged into one storage system but not another displays in the list, but</p>
--

will show up as logged out in the other storage system if they are added to the host.
If an Initiator is already in an host group on ALL of the storage systems where that initiator is logged in, then this initiator is filtered out of the Available list. Host groups with the same name and the same contents will be filtered to only show once.
If an initiator is not in an host group on one storage system, but it is in a host group on another storage system, both the initiator and the host group will be shown in the list.
Host groups with the same name but different contents will be shown individually with “Sym” and the last three digits of the storage appended to the name.
Host groups with different names but same contents across different storage systems display individually.
If an initiator that is not in a host group has the same name as a host group on a different storage system, then the host group is appended with (Group).
Cascaded host groups are filtered out.

## Select Storage Resource Pool

Use this dialog box to select a storage resource pool for the operation. Selecting **None** will remove the storage group from FAST control.

## Provisioning storage

This section describes how to make storage available to hosts:

## Creating host groups

### Before you begin

To perform this operation, you must be a StorageAdmin.

The storage system must be running Enginuity version 5876, or HYPERMAX OS 5977 or higher.

The maximum number of hosts allowed in a host group depends on the HYPERMAX OS:

For Enginuity 5876, the maximum allowed is 32.

For HYPERMAX OS 5977 or higher, the maximum allowed is 64.

This procedure explains how to create a host group (collection of hosts). For instructions on creating a host, refer to [Creating hosts](#) on page 292.

To create host groups:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3. Click **Create > Create Host Group**.
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 selected by default.
6. Optional: Do one of the following:
  - To create new hosts to add to the group, click **Create New Host**. For instructions on creating hosts, refer to [Creating hosts](#) on page 292.
  - To add existing hosts to the group:
    - 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.
    - Click **Add**.

Repeat these steps for each additional host.
  - To set the host port attributes:
    - Click **Set Host Group Flags** to open the **Set Host/Host Group Flags** dialog box.
    - Optional: Select a host whose flag settings you want to copy.

- Modify any of the [attributes](#), by selecting the corresponding Override option (thereby activating the Enable option) and enable (select) or disable (clear) the flag.
  - 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**.
  8. Do either of the following:
    - Click **Run Now** to start the task now.
    - 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 920 and [Previewing jobs](#) on page 920.


## Adding hosts to host groups

### Before you begin

To perform this operation, you must be a StorageAdmin.

To add hosts to host groups:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3.  Select the host group (or empty host) and click .
4. Click on the link in the **Hosts** field.
5. Click **Add Hosts**.
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 **Run Now** or **Add To Job List**.

## Removing hosts from host groups


### Before you begin

To perform this operation, you must be a StorageAdmin.

The storage system must be running Enginuity version 5876 or higher.

To add hosts to host groups:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3.  Select the host group (or empty host) and click .

4. Click on the link in the **Hosts** field.
5. Select the host and click **Remove**.
6. Click **OK**.

## Modifying host groups



### Before you begin

To perform this operation, you must be a StorageAdmin.

The storage system must be running Enginuity version 5876 or higher.

Modifying host groups:

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3. Do one of the following:
  - To modify a host group:
    - Select the host group and click **Modify**.
    - 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.
  - To add a Host:
    - Optional: To create a new host to add to the group, click **Create**. For help, refer to [Creating hosts](#) on page 292.
    - Select the host and click **Add**. To filter the list, type part of the host name. Repeat this step for each additional host.
  - To remove a Host:
    -  In the list of hosts, select the host, click  and then click **Delete**.
    - Click **OK**.
4. Repeat these steps for each additional host.
5. Do either of the following:
  - Click **Run Now** to start the task now.
  - 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 920 and [Previewing jobs](#) on page 920.

## Renaming hosts/host groups

### Before you begin

To perform this operation, you must be a StorageAdmin.

The storage system must be running Enginuity version 5876 or higher.

To rename host/host groups:



**Procedure**

1. Select the storage system.
2. Select **Hosts > Hosts**.
3. Select the host/host group and click **Modify**.
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 or host group port flags

To set host or host group port flags:

**Procedure**

1. Select the storage system.
2. Select **Hosts > Hosts**.
3.
 



Select a host, click  and then click **Set Flags** to open the **Set Host/Host Group Flags** dialog.
4. Optional: Select a host/host group whose flag settings you want to copy from the **Copy Flags from Other Host/Host Group** drop-down menu.
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 5876 or higher.

To delete hosts/host groups:

**Procedure**

1. Select the storage system.
2. Select **Hosts > Hosts**.
- 3.

Select the host/host group from the list and click

4. Click **Delete**, then click **OK** to confirm.

## Viewing hosts/host groups

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts** to open the Hosts list view.
3. 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.

Click  to view the host/host group details.

The following controls are available:

**Create Host**—[Creating hosts](#) on page 292

**Create Host Group**—[Creating host groups](#) on page 302

**Provision Storage to Host**—[Using the Provision Storage wizard](#) on page 100 or [Using the Provision Storage wizard](#) on page 108


**Modify**—[Modifying hosts](#) on page 294 or [Modifying host groups](#) on page 304

**Set Flags**—[Setting host or host group port flags](#) on page 296

**Delete**—[Deleting hosts/host groups](#) on page 296

## Viewing host/host group details

### Procedure

1. Select the storage system.
2. Select **Hosts > Hosts**.
3. Select the host/host group.
4. Click  to view the host/host group details.

---

**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.

**Hosts**—Number of hosts in the group. This field only displays for host groups.

**Masking Views**—Number of masking views with which the host/host group is associated.

**Initiators**—Number of initiators in the host/host group. For host groups, the value includes initiators in any child host groups.

**Host Groups**—Number of host groups in which this host is a member. This field only displays for individual hosts. .

**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.

**PowerPath Hosts**—Number of PowerPath hosts.

## Creating masking views

### Before you begin

The following explains how to mask volumes on storage systems running Enginuity 5876 or higher.

To create a masking view, you need to have created initiator groups, port groups, and storage groups. For instructions, refer to [Creating port groups](#) on page 316.

### Procedure

1. Select the storage system.
2. Select **Hosts > Masking view** to open the **Masking view** list view.
3. Click **Create** to open the **Create Masking View** dialog box.

4. Type the **Masking View Name**.

Masking view names must be unique from other masking views on the array and cannot exceed 64 characters. Only alphanumeric characters, underscores ( \_ ), and (-) are allowed. Masking view names are case-insensitive.

5. Select the **Host**.

6. Select the **Port Group**.

7. Select the **Storage Group**.

8. Optional: Manually set the host LUN addresses:

- a. Click **Set Dynamic LUNs** to open the **Set Dynamic LUNs** 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 LUNs** dialog box.

9. Click **OK**.

## Renaming masking views

### Procedure

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 **Rename**.
4. Type the new **Name**, and click **OK**.


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 the **File Masking Views** page (**System > System Dashboard > File Dashboard > File Masking Views**).

### Procedure

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  and then click **Delete** to open the **Delete Masking View** confirmation dialog box.

4. To unmap volumes in the masking view from their mapped ports, select **Delete Storage Group(s)**.
  5. Click **OK**.

## Viewing masking views


### Procedure

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.

To view a masking group's details, select it and click . The following properties are displayed:

- **Name**— User-defined masking view name.
- **Capacity (GB)**— Total capacity, in GB, of all volumes in the masking view.
- **Host**— Name of the associated host.
- **Port Group**— Name of the associated port group.
- **Storage Group**— Name of the associated storage group.
- **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.
- **Ports**— Number of ports contained in the masking view.
- **Volumes**— Number of volumes in the storage group contained in the masking view.

Depending on the options chosen, some of the following controls are available:

- **Create**— [Creating masking views](#) on page 307
- **Rename**— [Renaming masking views](#) on page 308
- **View Path Details**— [Viewing masking view connections](#) on page 309
- **Delete**— [Deleting masking views](#) on page 308

## Viewing 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 the **File Masking Views** page (**System > System Dashboard > File Dashboard > File Masking Views**).

### Procedure

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 Path Details** to open the masking view connections view.
4. 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.

#### 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

### Procedure

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 .

The following properties display:

- **Name**—Name of the masking view.
- **Host**—Name of the host.
- **Port group**—Name of the port group.
- **Storage group**—Name of the storage group.
- **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.
- **Ports**—Number of ports contained in the masking view.
- **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.

## Set Dynamic LUN Addresses

Use this dialog box to manually assign host LUN addresses for a masking operation.

### Procedure

1. Select the storage system.
2. Select **Hosts > Masking Views**.
3. Select a masking view and click **Create** to open the **Create Masking View** dialog box.
4. Click **Set Dynamic LUNs** to open the **Set Dynamic LUNs** dialog box.

This dialog box contains the following elements:

**Starting LUN**—LUN address assigned to the first volume.

**Apply Starting LUN**—Sets the address for the volume and keeps the dialog box open for additional operations.

**Next Available LUN**—Increments the Starting LUN address to the next available.

**Volumes to be masked**— Select the volumes you want to mask from the volumes list.

Enter the necessary values and click **OK**.

## Setting initiator port flags

**Procedure**

1. Select the storage system.
2. Select **Hosts > Initiators**.
3. Select an initiator and click **Set Host Flags** to open the **Set Initiator Flags** dialog.
4. Optional: Copy the attributes of an existing flag by selecting a flag under the **Copy Flags** drop-down menu.
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

**Before you begin**

Any changes made to an initiator's attributes affect the initiator and all its ports.

To set initiator attributes:

**Procedure**

1. Select the storage system.
2. Select **Hosts > Initiators**.
3. Select an initiator and click **Set Attributes**.

The initiator director: port, initiator, and optional alias names display.

4. Type the FCID (Fibre Channel ID) Value.
5. 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:

**Procedure**

1. Select the storage system.
2. Select **Hosts > Initiators**.
- 3.

Select an initiator, click  and then click **Rename Alias**.

4. Type a Node Name and Port Name.

On storage systems running Enginuity 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:

### Procedure

1. Select the storage system.
2. Select **Hosts > Initiators**.
- 3.

Select the initiator, click  and then click **Replace Initiator**.

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

### Procedure

1. Select the storage system.
2. Select **Hosts > Initiators**.
- 3.

Select the initiator, click  and select **Remove Masking Entry** to open the **Remove Masking Entry** dialog box.

4. Select the director and port.
5. Click **OK**.

## Viewing initiators

### Procedure

1. Select the storage system.
2. Select **Hosts > Initiators**.
3. Use the **Initiators** list view to view and manage initiators.

The properties and controls displayed in the view vary depending on the Engenuity 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.

**Hosts** — Number of hosts the initiator is associated with

**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 storage systems running Enginuity 5876 or higher.

To view the initiator's details, click 

The following controls are available:

**Set Host Flags**—[Setting initiator port flags](#) on page 311

**Set Attributes**—[Setting initiator attributes](#) on page 312


**Rename Alias**—[Renaming initiator aliases](#) on page 312

**Replace Initiator**—[Replacing initiators](#) on page 313

**Removing Masking Entry**—[Removing masking entries](#) on page 313

## Viewing initiator details

### Procedure

1. Select the storage system.
2. Select **Hosts > Initiators**.
3. Select the initiator from the list and click 
4. The following properties are displayed:

---

### Note

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.

**Hosts**—Number of hosts.

**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 5876 or higher.

**Masking Views**—Number of associated masking views, including the masking views that are associated with any cascaded relationships. This field only applies/appears for storage systems running Enginuity 5876 or higher.

**Volumes**—Number of volumes.

**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.

The following controls are available:

**Set Attributes**—[Setting initiator attributes](#) on page 312

**Set Host Flags**—[Setting initiator port flags](#) on page 311


**Rename Alias**—[Renaming initiator aliases](#) on page 312

**Replace Initiator**—[Replacing initiators](#) on page 313

**Removing Masking Entry**—[Removing masking entries](#) on page 313

## Viewing volumes associated with host initiator

### Procedure

1. Select the storage system.
2. Select **Hosts > Initiators**.
3. Select the initiator from the list and click .
4. Click on the number in the **Volumes** field.
5. Use this view to view and manage volumes associated with the initiator.

The following properties display:

**Name**—Volume name.

**Type**—Type of volume.

**Allocated %**—% of space allocated.

**Capacity (GB)**—Volume capacity in GBs.

**Status**—Volume status.

**Emulation**—Volume emulation.

**SRDF Group**—SRDF group the volume belongs to.

**Host Paths**—Host paths for the volume.

To see more volume properties, select the volume and click .

The following controls are available, depending on the Enginuity version running on the storage system:

**Create**—[Creating volumes](#) on page 178

**Expand**— [Expanding existing volumes](#) on page 191

**Delete**  [Deleting volumes](#) on page 188—

**Create SG**

- —HYPERMAX OS 5977 or later: [Creating storage groups](#) on page 112

**Set Volume Attributes**—[Setting volume attributes](#) on page 195

**Set Volume Identifiers**—[Setting volume identifiers](#) on page 196

**Set Volume Status**—[Setting volume status](#) on page 194

**Change Volume Configuration**—[Changing volume configuration](#) on page 190

**Replication QoS**—[QOS for replication](#) on page 197

**Duplicate Volume**—[Duplicating volumes](#) on page 188

**Expand Volume**— [Expanding existing volumes](#) on page 191

**Start Allocate/Free/Reclaim**—[Managing thin pool allocations](#) on page 244

**Stop Allocate/Free/Reclaim**— [Managing thin pool allocations](#) on page 244

**Map**—[Mapping volumes](#) on page 192



**Unmap**—[Unmapping volumes](#) on page 193

**Set SRDF GCM**—[Setting the SRDF GCM flag](#) on page 434

**Reset SRDF/Metro Identity**— [Resetting original device identity](#) on page 432

## Viewing details of a volume associated with initiator

### Procedure

1. Select the storage system.
2. Select **Hosts > Initiators**.
3. Select the initiator from the list and click .
4. Click on the number in the **Volumes** field.
5. Select the volume from the list and click  to see its details.

The following controls are available:

**Create**—To select the type of volume to create refer to [Creating volumes](#) on page 178.

## 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:

#### Procedure

1. Select the storage system.
2. Select **Hosts > Port Groups**.
3. Click **Create**.
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 the 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. Select **Run Now** or **Add To Job List**.

## Deleting port groups

#### Procedure

1. Select the storage system.
2. Select **Hosts > Port Groups**.
- 3.

Select the port group, click  and then 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:

#### Procedure

1. Select the storage system.
2. Select **Hosts > Port Groups**.
3. Select the port group and click .
4. Click on the number in the **Ports** field.
5. Click **Add Ports**.

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 Groups**—Number of port groups where the port is a member.

**Masking Views**—Number of associated masking views.

**Mapped Volumes**—Number of associated mapped volumes.

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.

To remove ports from port groups:

#### Procedure

1. Select the storage system.
2. Select **Hosts > Port Groups**.
3. Select the port group and click .
4. Click on the number in the **Ports** field.
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  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

To rename port groups:

### Procedure

1. Select the storage system.
2. Select **Hosts > Port Groups**.
3. Select the port group and click **Modify**.
4. Type the new port group Name and click **Apply**.

## Viewing port groups

### Procedure

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.

**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.

To view more details of a port group, select it and click .

The following controls are available:


**Create**—[Creating port groups](#) on page 316

**Modify**—[Renaming port groups](#) on page 319

**Delete**—[Deleting port groups](#) on page 317

## Viewing port groups details

### Procedure

1. Select the storage system.
2. Select **Hosts > Port Groups**.
3. Select the port group and click .
4. Use the port groups Details view to view and manage a port group.

The following properties display:

**Name**—User-defined port group name.

**Ports**—Number of ports in the group. Click on the number for more details

**Masking Views**—Number of masking views where the port group is associated. Click on the number for more details.

**Last Update**—Timestamp of the most recent changes to the port group.

**Host I/O (IO/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.

**Port 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**—[Creating port groups](#) on page 316

**Modify**—[Renaming port groups](#) on page 319

**Delete**  —[Deleting port groups](#) on page 317

## Viewing ports in port group

### Procedure

1. Select the storage system.
2. Select **Hosts > Port Groups**.
3. Select the port group and click .
4. Click on the number in the **Number of Ports** field.
5. 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**—[Adding ports to port groups](#) on page 317



**Remove**  —[Removing ports from port groups](#) on page 318

 —[Viewing port details](#) on page 321



## Viewing port details

### Procedure

1. Select the storage system.
2. Select **Hosts > Port Groups**.
3. Select the port group and click .
4. Click on the number in the **Ports** field.
5. Select a port and click .
6. Use the port Details view to view and manage a port.

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.

**Ports 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.

## Volume Set Addressing

An addressing scheme that uses virtual busses, targets, and LUNs to increase greatly the number of LUNs that can be addressed on a target port. Volume Set Addressing is supported for HP-UX.

## Viewing host IO limits

### Procedure

1. Select the storage system.
2. Select **Hosts > Port Groups**.
3. Select the port group and click .
4. Click on the link in the **Host I/O (IO/Sec)** or **Host I/O (MB/Sec)** fields.

The following properties display:

**Storage Group**—Storage group on which the limit is set.

**Quota State**—Whether the limit is set directly on the storage group (Defined) or through a cascaded relationship (Shared).

**Dir:Port**—Storage system director and port in the port group.

**Host I/O Limit (MB/Sec)**—Total host I/O limit on the listed port in MB/Sec. This value is the associated port group's I/O limit divided across its ports.

**Host I/O Limit (IO/Sec)**—Total host I/O limit on the listed port in IO/Sec. This value is the associated port group's I/O limit divided across its ports.

**Child Host I/O Limit (MB/Sec)**—Total child host I/O limit on the listed port in MB/Sec. This value is the associated port group's I/O limit divided across its ports.

**Child Host I/O Limit (IO/Sec)**—Total child host I/O limit on the listed port in IO/Sec. This value is the associated port group's I/O limit divided across its ports.

## Managing storage for Mainframe

The Mainframe Dashboard provides you with a single place to monitor and manage configured splits, CU images, and CKD volumes. To access the Mainframe Dashboard:

### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.

The Mainframe Dashboard is organized into the following panels:

- **CKD Compliance**
- **CKD Storage Groups**
- **Actions**
- **Summary**

### CKD Compliance panel

Displays how well CKD storage groups are complying with their respective service level policies, if applicable. All of the storage groups on the Mainframe Dashboard are organized into the following categories:

#### Total

All Mainframe storage groups on the array.



#### Stable

Number of storage groups performing within the service level targets.



indicates that there are no storage groups performing within the service level targets.

**Marginal**

Number of storage groups performing below service level targets. indicates that there are no storage groups performing below service level targets.

**Critical**

Number of storage groups performing well below service level targets.



indicates that there are no storage groups performing well below service level targets.

**No Service Level**

No service level compliance information.

**CKD Storage Groups panel**

Displays all of the Mainframe storage groups on the array. Double-click on a storage group to see more details as well as information on its compliance and volumes.

**Actions panel**

Displays the following links:

**Provision Storage**

Opens the Mainframe Provision wizard, which guides you through the process of provisioning storage for a mainframe. For more information, see [Using the Provision Storage wizard for mainframe](#) on page 104.

**Create CKD Volumes**

Opens the **Create Volume** dialog, from where you can create a CKD volume. For more information, see [Creating CKD volumes](#) on page 330.

**Summary panel**

Displays the following mainframe summary information:

**Splits**

The number of splits on the selected array. To view the list of splits, click **Splits**. For more information about viewing splits, see [Viewing splits](#) on page 327.

**CU Images**

The number of CU images on the selected array. To view the list of CU images, click **CU Images**. For more information about viewing CU images, see [Viewing CU images](#) on page 328.

### CKD Volumes

The number of CKD volumes on the selected array. To view the list of CKD volumes, click **CKD Volumes**. For more information about viewing CKD volumes, see [Managing volumes](#) on page 177.

## Provisioning storage for mainframe

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

Service level provisioning eliminates the need for storage administrators to manually assign physical resources to their applications. Instead, storage administrators specify the service level and capacity required for the application and the system provisions the storage group appropriately.

You can provision CKD storage to a mainframe host using the Provision Storage wizard. For specific instructions about how to provision storage for mainframe, refer to [Using the Provision Storage wizard for mainframe](#) on page 104.

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

To provision storage for Open Systems, refer to [Using the Provision Storage wizard](#) on page 100.

### Mapping CKD devices to CU images

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

For more information about how to map CKD devices to CU images, see the following tasks:

- [z/OS map from the CU image list view](#) on page 332
- [z/OS map from the volume list view](#) on page 333

## Using the Provision Storage wizard for mainframe

### Before you begin

- The storage system must be running HYPERMAX OS 5977 Q1 2016, or higher, and have at least one FICON director configured.
- Depending on the type of configuration selected, not all of the steps listed below might be required.

To provision storage to mainframe:

### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.
3. In the Actions panel, click **Provision Storage**. The Provision Storage wizard for mainframe is displayed.
4. In the **Create Storage Group** page, 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.

If you want to create an empty storage group, proceed to the final step after typing the storage group name.

5. 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.

6. Select an **Emulation** type. Available values are **CKD-3390** and **CKD-3380**.

7. 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.

Available values are:

Service level	Performance level	Use case
Diamond	Ultra high	HPC, latency sensitive
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.

For all-flash storage systems, the only service level available is Diamond and it is selected by default.

8. Type the number of **Volumes** and select either a **Model** or **Volume Capacity**.

Selecting a **Model** type automatically updates the **Volume Capacity** value. Alternatively, you can type the **Volume Capacity**.

---

**Note**

The maximum CKD volume size supported is 1182006 cylinders or 935.66 GB. It is possible to create an empty Storage Group with no volumes.

---

9. (Optional) Configure volume options:

---

**Note**

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

---

a.

Hover the cursor on the service level and click  .

b. Edit the **Volume Identifier**.

The following options are available:

**None**

Do not set a volume identifier.

**Name Only**

All volumes will have the same name. Type the name in the **Name** field.

**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. Type the name in the **Name** field.

**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. Type the name in the **Name** field.

- c. To **Allocate capacity for each volume** you are adding to the storage group, select this option. You can use this option only for newly created volumes, not existing volumes.
  - d. 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.
  - e. Click **OK**.
10. (Optional) To add a child storage group, do one of the following:
- On all-flash storage systems, click **Add Storage Group**.
  - On all other storage systems click **Add Service Level**.

Specify a **Name**, **Service Level**, **Volumes**, and **Model/Volume Capacity**.

Repeat this step for each additional child storage group. The maximum number of child storage groups allowed is 64.

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:
- 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 920 and [Previewing jobs](#) on page 920
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
12. On the **CU Image** page, select whether to use a **New** or an **Existing** CU image, and then do the following depending on your selection:
- **New:**
    - a. Specify the following information for the new CU image:
      - **CU Image Number**
      - **SSID**
      - **Base Address**
    - b. Select a **Split** with which to associate the CU image.
  - **Existing:**

- a. Select a CU image.
  - b. To specify a new value for the base address, click **Set Base Address**. For more information about setting the base address, refer to [Setting the base address](#) on page 337.
13. Click **Next**.
14. On the **Review** page, review the summary information displayed.  
 If the storage system is registered for performance, you can subscribe for compliance alerts for the storage group and run a suitability check to ensure that the load being created is appropriate for the storage system.  
 To enable compliance alerts, select **Enable Compliance Alerts**.  
 To run a suitability check, click **Run Suitability Check**.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing splits

### Before you begin

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

To view the splits list view:

### Procedure

1. Select the storage system.
2. Select **HOSTS > Mainframe**.
3. Click **Splits** to display the Splits list view.

The following properties are displayed:

#### Split Name

The user-defined name for the split.

#### Alpha Serial #

The alpha serial number of the split.

#### PAV State


Indicates what type of PAV is enabled on the split. The types are: HyperPAV, DynamicPAV, or SuperPAV (5978 only).

#### CU Images

The number of CU images associated with the split.

#### Ports

The number of FICON ports assigned to the split.

4. Select the split and click .

The following properties display:

**Split Name**

The user-defined name for the split.

**Alpha Serial #**

The alpha serial number of the split.

**PAV State**

Indicates if PAV is enabled on the split.

**Number of CU Images**

The number of CU images associated with the split.

**Number of Ports**

The number of FICON ports assigned to the split.

## Viewing CU images

To view the CU images list view:

**Procedure**

1. Select the storage system.
2. Click **Hosts > CU Images** to display the **CU Images** list view.
3. The following properties display:

**CU Image Number**

The CU image number.

**SSID**

The netmask prefix value of the IP interface.

**Split**

The name of the split containing the CU image.

**Number of Volumes**

The number of volumes mapped to the CU image.

**Storage Groups**

The number of storage groups containing volumes mapped to the CU image.

**Total Number of Base Addresses**


The total number of the base addresses configured on the CU image. The total includes used plus unused base addresses.

**Number of Aliases**

The number of aliases in use on the CU image.

**Status**

The status of volumes in the CU image.

To view more details, click 

The following controls are available:



- **z/OS Map** — [z/OS map from the CU image list view](#) on page 332
- **z/OS Unmap** — [z/OS unmap from the CU image list view](#) on page 333
- **Assign Alias Range** — [Adding an alias range to a CU image](#) on page 336
- **Remove Alias Range** — [Removing an alias range from a CU image](#) on page 337

## Viewing CU image details

To view the CU images detailed view:

### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.
3. Click **CU Images** to display the **CU Images** list view.
- 4.

Select the CU image and click .

The following properties display:

#### **CU Image Number**

The CU image number.

#### **SSID**

The CU SSID.

#### **Split**

The name of the containing split.

#### **Number of Volumes**

The number of volumes.

#### **Storage Groups**

The number of storage groups.

#### **Status**

The current status of the CU image.

#### **Total Number of Base Addresses**

The total number of base addresses configured on the CU image. The total includes used plus unused base addresses.

#### **Number of Available Base Addresses**

The number of available base addresses, in hexadecimal.

#### **Available Base Addresses**

The available base address ranges on the CU image.

#### **Next Available Base Address**

The next available base address, in hexadecimal.

#### **Number of Aliases**

The number of alias addresses.

#### **Alias Address Range**

The assigned alias address range, if applicable.

### PAV Aliasing

The type of PAV aliasing: HyperPAV, DynamicPAV, or SuperPAV (5978 only).

## Creating CKD volumes


### Before you begin

- The storage system must be running HYPERMAX OS 5977.810.784, or later, and have at least one FICON director configured.
- Depending on the type of configuration selected, not all of the steps listed below might be required.

### Procedure

1. Select the storage system.
2. Do one of the following:
  - Select **Storage > Volumes**. In the **Volumes** list view, click **Create**.
  - Select **Hosts > Mainframe**. In the **Actions** panel, click **Create CKD volumes**.
3. Select the **Configuration** type.
4. From the **Emulation** list, select one of the following values:
  - **CKD-3390**
  - **CKD-3380**

5. Specify the capacity by typing the **Number of Volumes**, and selecting a **Volume Capacity**.

If the **Model** menu is available, selecting a model automatically updates the volume capacity to the correct capacity. Alternatively, you can manually enter a volume capacity by clicking  .

6. (Optional) To add the volumes to a CKD storage group, click in the **Add to Storage Group** field to reveal a drop-down menu of available CKD storage groups. Click **Clear** to clear the selection.
7. Click **Advanced Options**

The advanced options that are presented depend on the configuration details. Complete any of the following steps that are appropriate:

- a. If required, type an **SSID** or click **Select** to choose one.
- b. To name the new volumes, select one of the following Volume Identifiers:

#### 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 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 196.

- c. Depending on the value selected for **Volume Identifier**, type a **Name**, or a **Name and Append Number**.
  - d. 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.
  - e. Click **OK**.
8. Do one of the following:
- a. 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 920 and [Previewing jobs](#) on page 920.
  - b. Expand **Add to Job List** and click **Run Now** to perform the operation now.

## Editing CKD volume capacities

### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe > Create**.
3. Click  to open the **Edit Volume Capacities** dialog.
4. Use the drop-down menus to choose the number of volumes, the model, the capacity and the unit to be used to measure capacity (TB, GB, MB or cylinders).  
Click  to add another volume size.
5. Click **Apply** to apply your changes or **Cancel** to reject them.

## Expanding CKD volumes

### Before you begin

To expand CKD volumes requires HYPERMAX OS 5977.1125.1125 or later. In addition, you must be logged in as an Administrator.

You can expand a volume up to 1,182,006 cylinders (1 TB). When expanding a device above 565,250 cylinders, the new size must be a multiple of 1113 cylinders. If you specify a size that isn't that multiple, the system rounds the size up to the next multiple of 1113.

You cannot expand a volume when it is:

- A CKD 3380 device
- A TDAT

- Marked as Soft Fenced
- Part of a RDF session
- Part of a SnapVx session

The procedure below shows one way to expand a CKD volume. You can also carry out this task via **Storage > Storage Groups > Volumes** or **Storage > Volumes**.

#### Procedure

1. Select the storage system.
2. Click **Hosts > Mainframe**, click **CKD Volumes**, select a volume and click **Expand** to open the **Expand Volume** dialog.
3. 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.
4. To reserve the volume, select **Reserve Volumes**.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## z/OS map from the CU image list view

#### Before you begin

The storage system must be running HYPERMAX OS 5977 Q1 2016.

#### Note

Before making any mapping changes to an existing CU image, please ensure that all of the devices in the CU are offline (the status of the CU should be offline).

To map to a CU image from the CU image list view:

#### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.
3. Click **CU Images**.
4. Select a CU image, which has not already been mapped, and click **z/OS Map**.  
The CU Image Map wizard displays.
5. In the **Find Volumes** page, search for a volume to which you can map the CU image:
  - a. (Optional) Specify one or more criteria by which you can filter volumes.  
An **Additional Criteria** filter for volumes with emulation CKD-3390 is applied by default.
  - b. (Optional) Click **Add Another** to configure further additional criteria.
  - c. Click **Find Volumes**.
6. In the **Select Volumes** page, select one or more volumes to map to the CU image.

7. Click **Summary**.
8. Review the summary information.
9. (Optional) To reset the base address, click **Set Base Address** and specify the new base address.
10. Do one of the following:
  - a. 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 920 and [Previewing jobs](#) on page 920.
  - b. Expand **Add to Job List** and click **Run Now** to perform the operation now.

## z/OS unmap from the CU image list view

### Before you begin

The storage system must be running HYPERMAX OS 5977 Q1 2016.

---

### Note

Before making any mapping changes to an existing CU image, please ensure that all of the devices in the CU are offline (the status of the CU should be offline).

---

To unmap a CU image from the CU image list view:

### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.
3. Click **CU Images**.
4. Select the CU image you want to unmap. Click **z/OS Unmap**.  
The **CU Image Unmap** dialog box displays.
5. Select one or more volumes to unmap from the CU image.
6. Do one of the following:
  - a. 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 920 and [Previewing jobs](#) on page 920.
  - b. Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## z/OS map from the volume list view

### Before you begin

The storage system must be running HYPERMAX OS 5977 Q1 2016 or higher.

---

### Note

Before making any mapping changes to an existing CU image, please ensure that all of the devices in the CU are offline (the status of the CU should be offline).


---


To map to a CU image from the volume list view:

### Procedure

1. Select the storage system.

2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.
3. Click **CKD Volumes**.
4.
 



Select one or more volumes to map, click  and then click **z/OS Map**.

The **Mainframe Volumes Mapping** dialog box displays.
5. Select whether to want to map the volume(s) to a **New** or an **Existing** CU image.
  - **New**
    - a. Specify values for **CU Image Number**, **SSID**, and **Base Address**.
    - b. (Optional) Select a **Split**.
  - **Existing**
    - a. Select the CU image to which you want to map the selected volume(s).
    - b. (Optional) Click **Set Base Address** to reset the next available base address.
6. Do one of the following:
  - a. 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 920 and [Previewing jobs](#) on page 920.
  - b. Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## z/OS unmap from the volume list view

### Before you begin

The storage system must be running HYPERMAX OS 5977 Q1 2016 or higher.


### Note


Before making any mapping changes to an existing CU image, please ensure that all of the devices in the CU are offline (the status of the CU should be offline).

To unmap a CU image (from the volume list view):

### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.
3. Click **CKD Volumes**.
4.
 



Select one or more volumes to unmap, click  and then click **z/OS Unmap**.

The **Mainframe Volumes Unmapping** dialog box displays a summary of the unmap operation.
5. Do one of the following:
  - a. 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 920 and [Previewing jobs](#) on page 920.
  - b. Expand **Add to Job List** and click **Run Now** to perform the operation now.

## z/OS map from the Volumes (Storage Groups) list view

### Before you begin



The storage system must be running HYPERMAX OS 5977 Q1 2016 or higher.

### Note

Before making any mapping changes to an existing CU image, please ensure that all of the devices in the CU are offline (the status of the CU should be offline).

To map to a CU image from the Volumes (Storage Groups) list view:

### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe**.
3. In the **CKD Storage Groups** panel, click **View All Storage Groups** to open the **Storage Groups** list view.
4. Select the storage group and click  to see its details. Click on the number in the **Number of Volumes** field to open the **Volumes (Storage Groups)** list view.
5. Select one or more volumes to map, click  and then click **z/OS Map**.  
The **Mainframe Volumes Mapping** dialog box displays.
6. Select whether to want to map the volume(s) to a **New** or an **Existing** CU image.
  - **New**
    - a. Specify values for **CU Image Number**, **SSID**, and **Base Address**.
    - b. (Optional) Select a **Split**.
  - **Existing**
    - a. Select the CU image to which you want to map the selected volume(s).
    - b. (Optional) Click **Set Base Address** to reset the next available base address.
7. Do one of the following:
  - a. 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 920 and [Previewing jobs](#) on page 920.
  - b. Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## z/OS unmap from the Volumes (Storage Groups) list view

### Before you begin

The storage system must be running HYPERMAX OS 5977 Q1 2016 or higher.

---



**Note**

Before making any mapping changes to an existing CU image, please ensure that all of the devices in the CU are offline (the status of the CU should be offline).

---

To unmap a CU image (from the Volumes (Storage Groups) list view):

**Procedure**

1. Select the storage system.
2. Select **Hosts > Mainframe**.
3. In the **CKD Storage Groups** panel, click **View All Storage Groups** to open the **Storage Groups** list view.
4. Select the storage group and click  to see its details. Click on the number in the **Number of Volumes** field to open the **Volumes (Storage Groups)** list view.
5. Select one or more volumes to unmap, click  and then click **z/OS Unmap**.  
The **Mainframe Volumes Unmapping** dialog box displays a summary of the unmap operation.
6. Click **Yes** to the warning dialog box.  
The **Mainframe Volumes Unmapping** dialog box displays a summary of the unmap operation.
7. Do one of the following:
  - a. 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 920 and [Previewing jobs](#) on page 920.
  - b. Expand **Add to Job List** and click **Run Now** to perform the operation now.

## Adding an alias range to a CU image

**Before you begin**

The storage system must be running HYPERMAX OS 5977 Q1 2016.

To add an alias range to a CU image:

**Procedure**

1. Select the storage system.
2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.
3. Click **CU Images**, select the CU image to which you want to add an alias range and click **Assign Alias Range**.
4. Type the **Start Alias** (Next available address).  
The minimum value allowed is 00.
5. Type the **End Alias**.  
The maximum value allowed is FF.
6. If required, select **Reserve Volumes**.



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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List** and click **Run Now** to perform the operation now.

## Removing an alias range from a CU image

### Before you begin

- The storage system must be running HYPERMAX OS 5977 Q1 2016.
- This operation removes all of the aliases for the selected CU image.

To remove an alias range from a CU image:

### Procedure

1. Select the storage system.
2. Select **Hosts > Mainframe** to open the Mainframe Dashboard.
3. Click **CU Images**, select the CU image from which you want to remove an alias range and click **Remove Alias Range**.
4. Review the information displayed in the **Remove Alias Range** dialog box 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List** and click **Run Now** to perform the operation now.

## Setting the base address

The **Set Base Address** dialog box is launched from the following locations:

- Provision Storage wizard — [Using the Provision Storage wizard](#) on page 100
- CU Image Map wizard — [z/OS map from the CU image list view](#) on page 332
- **Mainframe Volumes Mapping** dialog — [z/OS map from the volume list view](#) on page 333

To set the base address:

### Procedure

1. In the **Base Address** field, specify a new value for base address.  
Addresses in the range 00-FF are allowed.
2. Click **OK**.

## Understanding All Flash Mixed FBA/CKD support

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

---

**Note**

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

- Purchased as a mixed All Flash system
  - Installed at HYPERMAX OS 5977 Q2 2017 or later
  - Configured with 2 Storage Resource Pools - 1 FBA Storage Resource Pool and 1 CKD Storage Resource Pool
- 

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

For specific instructions about how to provision storage for mainframe, refer to [Using the Provision Storage wizard for mainframe](#) on page 104, by default only the CKD SRP is available in the **Storage Resource Pool** drop down list.

To provision storage for Open Systems, refer to [Using the Provision Storage wizard](#) on page 100, by default only the FBA SRP is available in the **Storage Resource Pool** drop down list.

For specific instructions about how to modify a storage group, refer to [Modifying storage groups](#) on page 119, depending on the storage group selection the **Storage Resource Pool** drop down list is filtered to display the CKD or FBA SRP.

---

**Note**

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

**Mapping FBA devices to CU images**

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

For more information about how to map FBA devices to CU images, see the following tasks:

- [z/OS map FBA volumes from the Volumes \(Storage Groups\) list view \(HYPERMAX OS 5977 or higher\)](#) on page 338
- [z/OS unmap FBA volumes from the Volumes \(Storage Groups\) list view](#) on page 339

**z/OS map FBA volumes from the Volumes (Storage Groups) list view (HYPERMAX OS 5977 or higher)**

**Before you begin**

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

- Purchased as a mixed all flash system

- Installed at HYPERMAX OS 5977 Q2 2017 or later
- Configured with 2 Storage Resource Pools - 1 FBA Storage Resource Pool and 1 CKD Storage Resource Pool

See [Understanding All Flash Mixed FBA/CKD support](#) on page 337 for additional information.

---



#### Note

Before making any mapping changes to an existing CU image, please ensure that all of the devices in the CU are offline (the status of the CU should be offline).

---

To map to a CU image from the Volumes (Storage Groups) list view:

#### Procedure

1. Select the storage system.
2. Select **Storage > Storage Groups**.
3. Select a storage group and click  to see its details
4. Click on the number in the **Volumes** field to open the **Volumes (Storage Groups)** list view.
5. Select one or more volumes to map, click  and then click **z/OS Map**.  
The **Mainframe Volumes Mapping** dialog box displays.
6. Select whether to want to map the volume(s) to a **New** or an **Existing** CU image.
  - **New**
    - a. Specify values for **CU Image Number**, **SSID**, and **Base Address**.
    - b. (Optional) Select a **Split**.
  - **Existing**
    - a. Select the CU image to which you want to map the selected volume(s).
    - b. (Optional) Click **Set Base Address** to reset the next available base address.
7. Do one of the following:
  - a. 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 920 and [Previewing jobs](#) on page 920.
  - b. Expand **Add to Job List**, and click **Run Now** to perform the operation now.

### z/OS unmap FBA volumes from the Volumes (Storage Groups) list view

#### Before you begin

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

- Purchased as a mixed all flash system
- Installed at HYPERMAX OS 5977 Q2 2017 or later
- Configured with 2 Storage Resource Pools - 1 FBA Storage Resource Pool and 1 CKD Storage Resource Pool



See [Understanding All Flash Mixed FBA/CKD support](#) on page 337 for additional information.

#### Note

Before making any mapping changes to an existing CU image, please ensure that all of the devices in the CU are offline (the status of the CU should be offline).

To unmap a CU image (from the Volumes (Storage Groups) list view):

#### Procedure


1. Select the storage system.
2. Select **Storage > Storage Groups**.
3. Select a storage group and click  to see its details
4. Click on the number in the **Volumes** field to open the **Volumes (Storage Groups)** list view.
5. Select one or more volumes to unmap, click  and then click **z/OS Unmap**.  
The **Mainframe Volumes Unmapping** dialog box displays a summary of the unmap operation.
6. Do one of the following:
  - a. 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 920 and [Previewing jobs](#) on page 920.
  - b. Expand **Add to Job List** and click **Run Now** to perform the operation now.

## 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.

#### Procedure

1. Select the storage system.
2.
  - To map at the volume level:
    - a. Select **Storage > Volumes**.
    - b. To display only CKD volumes, click in the **Emulation** field and select CKD from the drop-down menu.
    - c. Select a CKD volume, click  and then click **z/OS Map** to open the **z/OS Map Volumes** dialog box.

---

**Note**

To create a new CU Image - Enter a base address with a 4 hexadecimal format e.g. "3210"

- "32" = the CU Image Id
- "10" = base address (First base address must end with 0)

To create a new SSID - Enter a SSID with a 4 hexadecimal format e.g. "1234" (must be unique)

---

- To map at the CU image level:
  - a. Select **Hosts > CU Images**.
  - 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 920 and [Previewing jobs](#) on page 920.
  - **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.

### Procedure

1. Select the storage system.
2.
  - To unmap at the volume level:
    - a. Select **Storage > Volumes**.
    - b. To display only CKD volumes, click in the **Emulation** field and select CKD from the drop-down menu.
    - c.
 



Select a CKD volume, click  and then 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Copying CU image mapping

### Before you begin

Before you begin:

All volumes in a specified range must be mapped to the same CU image, or not mapped at all. Volumes within the specified range that are not mapped will be ignored as long as they are not mappable (SAVE devices, DRVs, and so on). If a volume in the specified range is mappable, the request will be rejected.

The following explains how to copy the front-end mapping addresses of a set of volumes from one port to another, providing multi-path access from the storage system to the mainframe.

To copy CU image mapping:

### Procedure

1. Select **Hosts > CU Images** to open the **CU Images** list view.
2. Select an image, and click **Copy Mapping** to open the **z/OS Map** dialog box.

## Available Volume for EA/EF Mapping dialog box

Use this dialog box to select one or more volumes for the mapping operation.

To select a range of volumes, select the first volume in the range, press and hold the **Shift** key, and then click the last volume in the range.

## Base Addresses in Use dialog box

Use this dialog box to view base addresses already in use.

## Select SSID dialog box


Use this dialog box to select an SSID for the operation.

## Viewing CKD volumes in CU image

Viewing CKD volumes in CU image

### Procedure

1. Select the storage system.

2. Select **Hosts > CU Images**
3. Select the CU image and click .
4. In the details panel, click on the number in the **Number of Volumes** field to open the **CKD Volumes** list view..
5. Use the **CKD Volumes** list view to display and manage CKD volumes in a CU image.

### Results

**Name** — Symmetrix volume name.

**Type** — Volume configuration.

**Status** — Volume status.

**Capacity (GB)** — Volume capacity in GBs.

**Emulation** — Emulation type.

**UCB Address** — Unit control block (address used by z/OS to access this volume).

**Volser** — Volume serial number (disk label (VOL1) used when the volume was initialized).

The following controls are available:



— [Viewing CU image details](#) on page 329

**z/OS Map**—[z/OS map from the volume list view](#) on page 333

**z/OS Unmap**—[z/OS unmap from the volume list view](#) on page 334

## Creating PowerPath hosts

### Before you begin

The following are the minimum requirements to perform this task:

- A storage system running PowerMax 5978 or higher.
- Unisphere for PowerMax 9.0.
- Solutions Enabler 9.0.
- PowerPath 6.3.

### Procedure

1. Select the storage system.
2. Select **Hosts > PowerPath Hosts** to open the **PowerPath Hosts** list view.
3. Click **Create Host** to open the **Create Host for PowerPath Host** dialog.
4. You can use the host name that appears in the dialog or else type in a new one.  
Host names must be unique from other host/host group names on the storage system.
5. Select either **Add To Job List** or **Run Now**.

All initiators associated with the selected PowerPath Host will be added to the new host.

## Viewing PowerPath hosts

### Before you begin

The following are the minimum requirements to perform this task:

- A storage system running PowerMaxOS 5978 or higher.
- Unisphere for PowerMax 9.0.
- Solutions Enabler 9.0.
- PowerPath 6.3.

### Procedure

1. Select the storage system.
2. Select **Hosts > PowerPath Hosts** to open the **PowerPath Hosts** list view.

The following properties display:

- **Name** — The PowerPath host name.
- **Version** — The PowerPath host version.
- **OS Version** — The PowerPath host OS version.
- **Vendor** — The PowerPath host hardware vendor.
- **Initiators** — The number of PowerPath host initiators.
- **Hosts** — The number of PowerPath hosts.
- **VMs** — The number of PowerPath host virtual machines.

The following control is available: **Create Host:** [Creating PowerPath hosts](#) on page 343


## Viewing PowerPath hosts details

### Before you begin

The following are the minimum requirements to perform this task:

- A storage system running PowerMaxOS 5978 or higher.
- Unisphere for PowerMax 9.0.
- Solutions Enabler 9.0.
- PowerPath 6.3.

### Procedure

1. Select the storage system.
2. Select **Hosts > PowerPath Hosts** to open the **PowerPath Hosts** list view.
3. To view the details of a PowerPath host, select it and click 

The following properties display:


- **Name** — The PowerPath host name.
- **Version** — The PowerPath host version.



- **Patch Level** — The PowerPath host patch level.
- **License Info** — The PowerPath host license info.
- **Vendor** — The PowerPath host hardware vendor.
- **OS Version** — The PowerPath host OS version.
- **OS Revision** — The PowerPath host OS revision.
- **Host Registration Time** — The time the host registered with the POWERMAX array.
- **Connectivity Type** — Indicates whether the PowerPath Host is connected to the POWERMAX array by iSCSI or Fibre.
- **Cluster Name** — The PowerPath host cluster name.
- **Cluster Node Name** — The PowerPath host node name in the cluster.
- **Initiators** — The number of PowerPath host initiators. Click on the number to see the initiators list view.
- **Hosts** — The number of PowerPath hosts.
- **Masking Views** — The number of PowerPath host masking views. Click on the number to see the masking views list view.
- **VMs** — The number of PowerPath host virtual machines. Click on the number to see the VMs list view.
- **Storage Groups** — The number of PowerPath host storage groups. Click on the number to see the storage groups list view.
- **Volumes** — The number of PowerPath host volumes. Click on the number to see the volumes list view.

## Viewing PowerPath Host Virtual Machines

### Procedure

1. Select the storage system.
2. Select **Hosts > PowerPath Hosts** to open the **PowerPath Hosts** list view.
3. Select a PowerPath Host and click the information icon .
4. In the details panel, click on the number in the **VMs** field to open the **VMs** list view.

The following properties display:

- **Name** — The VM name.
- **OS Name** — The VM operating system.

A link allows you to add ESX or vCenter viClient credentials so as to retrieve more information on the Virtual Machine. In that case, the following additional properties are displayed:

- **Power State** — The VM power status.
- **CPU Count** — The number of CPUs assigned to the VM.
- **Total Memory** — The total RAM assigned to the VM.
- **State** — The current state of the VM.
- **Address** — The IP address of the VM.

## Viewing host cache adapters

### Procedure

1. Select the storage system.
2. Select **Hosts > Xtrem SW Cache Adapters** to open the **XtremSW Cache Adapters** list view.

The following properties display:

**Card S/N**—Adapter serial number.

**Card Version**—Adapter version.

**Vendor**—Adapter vendor.

**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

• Understanding Data Protection Management.....	348
• Creating device groups.....	348
• Understanding TimeFinder/Clone operations.....	354
• Understanding TimeFinder/Snap operations.....	367
• Managing TimeFinder/Mirror sessions.....	378
• Managing TimeFinder SnapVX.....	385
• Managing remote replication sessions.....	402
• Understanding Virtual Witness.....	444
• Creating SRDF/A DSE pools.....	448
• Creating TimeFinder/Snap pools.....	451
• Viewing SRDF group volumes.....	454
• Viewing SRDF protected storage groups.....	455
• Viewing related SRDF groups.....	457
• Creating SRDF groups.....	457
• Understanding RecoverPoint.....	471
• Creating Open Replicator copy sessions.....	490
• Understanding non-disruptive migration (NDM).....	500
• Viewing the authorized users and groups details.....	510
• Expanding remote volumes .....	511
• Setting a device identity.....	511
• Editing storage group volume details .....	512
• Editing storage group details .....	513
• Replication state severities.....	513
• Managing space reclamation.....	514
• Advanced Options dialog.....	515

# Understanding Data Protection Management

Data Protection Management covers the following areas:

- **Storage Groups** - Management of SRDF protected storage groups.
- **Device Groups** - Management of device groups. A device group is a user-defined group comprised of devices that belong to a locally attached array. Control operations can be performed on the group as a whole, or on the individual device pairs in the group. By default, a device can belong to more than one device group.
- **SRDF Groups** - Management of SRDF groups. SRDF groups provide a collective data transfer path linking volumes of two separate storage systems. These communication and transfer paths are used to synchronize data between the R1 and R2 volume pairs associated with the RDF group. At least one physical connection must exist between the two storage systems within the fabric topology.
- **Migrations** - Non-disruptive migration (NDM) management. NDM allows you to migrate storage group (application) data in a non-disruptive manner with no downtime from NDM capable source arrays to NDM capable target arrays.
- **SRDF/A DSE Pools** - Management of SRDF/A DSE Pools.
- **TimeFinder Snap pools** - Management of TimeFinder Snap pools.
- **Open Replicator** - Management of Open Replicator. Open Replicator (ORS) provides a method for copying data to or from various types of arrays within a storage area network (SAN) infrastructure.
- **RecoverPoint Systems** - Management of RecoverPoint systems.
- **Virtual Witness** - Management of Virtual Witness. The Witness feature supports a third party that the two storage systems consult if they lose connectivity with each other, that is, their SRDF links go out of service. When this happens, the Witness helps to determine, for each SRDF/Metro Session, which of the storage systems should remain active (volumes continue to be read and write to hosts) and which goes inactive (volumes not accessible).

## Creating device groups

Target volumes are automatically created by the wizard when the source Storage Group contains CKD volumes.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **Device Groups** tab 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.
- **RDF1** — Group can only contain R1 volumes.

- RDF2 — Group can only contain R2 volumes.
  - RDF21 — 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**.
    - 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**.
  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**.

### Results

A window appears that displays the progress of the wizard's tasks.

## Adding volumes to device groups


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.

3. Click the **Device Groups** tab to open the **Device Group** list view.
4. Select the device group and click **Add Volumes**.
5. From the list of available volumes, select the volume(s) and click **Add to Group**.
6. Optional: Remove a previously added volume by selecting it and clicking **Remove**.
7. Click **OK**.

## Removing volumes from device groups

### Procedure



1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **Device Groups** tab to open the **Device Group** list view.
4. Select the device group and click  to open the **Details** view.
5. Click the number next to **Number of Volumes** to view all volumes in device group.
6. Select one or more volumes and click **Remove Volumes**.
7. Click **OK**.

## Setting consistency protection

### Before you begin

To set consistency protection:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
3.   
Select a group, click more , and select **Asynchronous > Set Consistency**.
4. select **Enable** or **Disable**.
5. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
6. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
7. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.


## Renaming device groups

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **Device Groups** tab to open the **Device Group** list view.
4. Select the device group from the list and click **Rename**.
5. In the **Name** field, enter the new device group name.
6. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Deleting device groups

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **Device Groups** tab to open the **Device Group** list view.
4. Select the device group and click  .
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing device groups

### Procedure



1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **Device Groups** tab 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).
- **Group Valid**—Indicates whether the device group is valid.

The following controls are available, depending on the operating environment:

-  — [Viewing device group details](#) on page 352
- **Create**—[Creating device groups](#) on page 348
- **Rename**—[Adding volumes to device groups](#) on page 349
-  — [Deleting disk groups](#) on page 237
- **Add Volumes**—[Adding volumes to device groups](#) on page 349
- **Replication QOS**—[QOS for replication](#) on page 197
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945 (Does not apply/display with HYPERMAX OS 5977 or higher.)
- **Assign Symmetrix Priority**—[Assigning array priority to groups of volumes](#) on page 189 (Does not apply/display with HYPERMAX OS 5977 or higher.)
- **Set Optimized Read Miss**—[Setting optimized read miss](#) on page 193 (Does not apply/display with HYPERMAX OS 5977 or higher.)

## Viewing device group details

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **Device Groups** tab to open the **Device Group** list view.
4. Select the device group and click  to open the **Details** view.

The following properties display, depending on the operating environment:

- **Name**—User-defined device group name.
- **Application ID**—Indicates which application created the device group.
- **Group Valid**—Indicates whether the device group is valid.
- **Device Group Create Time**—Time the device group was created.
- **Device Group Modify Time**—Time the device group was modified.
- **Symmetrix ID**—Storage system serial number ID.
- **Number of Volumes**—Number of volumes.
- **Number of Associate 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.
- **Number of Composite Groups**—Number of composite groups.
- **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.

Links are provided 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 the number next to **Number of Volumes** opens the view listing the volumes contained in the device group.

## Viewing volumes in device group

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **Device Groups** tab to open the **Device Group** list view.

4. Select the device group and click  to open the **Details** view.
5. Click the number next to **Number of Volumes** to view all volumes in device group.

The following properties display:

- **Name**—Volume name
- **LDev**—Logical device name
- **Volume Config**—Device Configuration
- **Capacity (GB)**—Device capacity in GB
- **Status**—Device status

The following controls are available:

- **Add Volumes**—[Adding volumes to device groups](#) on page 349
- **Remove Volumes**—[Removing volumes from device groups](#) on page 350

## 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

### Before you begin

TimeFinder/Clone requires Engenuity version 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.

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


To manage TimeFinder/Clone sessions:

## Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.

The following properties display:

- **Device 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. Property values: RDF1, RDF2, RDF21, and Regular.
- **Group Valid**—Indicates whether the group is valid. Property values: Yes or No.

Click  and click the number next to **Clone Pairs** to view the associated clone pairs (see [Viewing clone pairs](#) on page 364).

Click  to click the number next **Storage Groups** to view the associated storage groups.

The following controls are available:

- **Create Pairs**—[Creating clone copy sessions](#) on page 355
- **Activate**—[Activating clone copy sessions](#) on page 357
- **Recreate**—[Recreating clone copy sessions](#) on page 358
- **Split**—[Splitting clone volume pairs](#) on page 362
- **Restore**—[Restoring data from target volumes](#) on page 361
- **Establish**—[Creating Snapshots](#) on page 379
- **Terminate**—[Terminating clone copy sessions](#) on page 363
- **Set Mode**—[Modifying clone copy sessions](#) on page 360

## Creating clone copy sessions

This procedure explains how to create clone copy sessions.

---

**Note**

Note the following:

- TimeFinder/Clone requires Engenuity version 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\_GRP 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.
- 

**Procedure**

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- a. Select a group and click **Create Pairs**.
- b. Select a source type and a target type.

Pair level:

- a. Select a group and click ⓘ to open its **Details** view.
  - b. Click on the number next to **Clone Pairs**.
  - c. Select one or more pairs and click **Create Pairs**.
  - d. Click **Set Pairs** to open the **Set Pairs** dialog box.
  - e. Select a source volume and a target volume, and click **Add** to make them a pair. Repeat this step as required.
  - f. Click **OK** to return to the **Create Sessions** dialog box.
5. Click **Advanced Options** to set the advanced options as described next.

Setting Advanced Options:

If performing this operation at the group level, you can optionally select a **Pairing Type** and select one of the following. If you are not using the **Pairing Type** option, leave this field set to **None**.

- **Use Exact Pairs**—Allows the system to pair up the volumes in the exact order that they were added to the group.

- **Use Optimized Pairs**—Optimizes volume pairings across the local Symmetrix system without regard for whether the volumes belong to different RDF (RA) groups.

For Copy Mode, 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 No Copy**—Specifies to change the session to CopyOnAccess once the session is activated and no full-volume copy will initiate.
- **Use PreCopy**—Specifies to start copying tracks in the background before you activate the clone session.

By default, when creating a clone session, the system creates 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**.

To attach Session Options to the operation, select any number of [options](#).

Click **OK**.

6. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - 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 a 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.

---

### Note

- TimeFinder/Clone requires Engenuity version 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 the copy operation from the source volume to the target volume:


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
4. Do one of 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**.
- Select a source type and a target type.

Pair level:

- a. Select a group and click  to open its **Details** view.
  - b. Click on the number next to **Clone Pairs**.
  - c. Select one or more pairs, and click **Activate**.
5. Optional: To attach session options to the operation, click **Advanced Options** 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Recreating clone copy sessions

### Before you begin

- TimeFinder/Clone requires Enginuity version 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.

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.

### Procedure


- To recreate clone copy sessions:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Device Groups**.
  3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
  4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

#### Group level:

Select a group and click **Recreate**.

Select a source type and a target type.

#### Pair level:

- a. Select a group, and click  to open its **Details** view.
  - b. Click on the number next to **Clone Pairs**.
  - c. Select one or more pairs and click **Recreate**.
5. Optional: To attach session options to the operation, click **Advanced Options**, 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Creating clone snapshots

### Before you begin


- TimeFinder/Clone requires Enginuity OS 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.

This procedure explains how to create and immediately activate clone snapshots



### Procedure

1. To create clone snapshots:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Device Groups**.
  3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
  4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

#### Group level:

- a.
  1. Select a group, click , and select **Create Snapshot**.
  2. Select the source type and the target type.

#### Pair level:

- a. Select a group and click  to open its **Details** view.
- b. Click on the number next to **Clone Pairs**.
- c.
  1. Select one or more pairs, click , and select **Create Snapshot**.
  2. Select the source type and the target type.
5. Specify whether to perform an **Incremental** or **Full** create.

6. Optional: To attach session options to the operation, click **Advanced Options**, 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Modifying clone copy sessions

### Before you begin


- TimeFinder/Clone requires Enginuity version 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 change a session created with the Differential option to the No Copy mode, as the session will fail.

This procedure explains how to modify the mode in which a clone copy session is operating.


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

#### Group level:

Select a group, click , and select **Set Mode**.

#### Pair level:

- a. Select a group, and click  to open its **Details** view.
- b. Click on the number next to **Clone Pairs**.
- c.

Select one or more pairs, click , and select **Set Mode**

5. 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.

6. If performing the operation at the group level, select the type of source volumes (**Source Type**) and the type of target volumes **Target Type**.
7. Optional: To set session options, click **Advanced Options**, and select any number of [options](#).
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Restoring data from target volumes

### Before you begin

- TimeFinder/Clone requires Enginuity version 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Clone operations are mapped to their TimeFinder/SnapVX equivalents.
- With Enginuity 5876 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 5876, 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.

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.

To restore data from a target volume:


### Procedure


1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

#### Group level:

Select a group, click , and select **Restore**.

**Pair level:**

- a. Select a group, and click  to open its **Details** view.
- b. Click on the number next to **Clone Pairs**.

- c.
  - Select one or more pairs, click , and select **Restore**.

5. Select a **Restore Type**:

- **Incremental** — 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.
- **Full** — Terminates the original session and starts a copy session to the target of the restore.

6. If performing the operation at the group level, select the type of source volumes (**Source Type**) and the type of target volumes **Target Type**.

7. To attach **Session Options** to the operation, click **Advanced Options**, and select any number of [Clone copy session options](#) on page 365.

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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now

## Splitting clone volume pairs

### Before you begin

- TimeFinder/Clone requires Enginuity version 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.

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.



### Procedure

1. To split clone volume pairs:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Device Groups**.
  3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
  4. Do one of the following, depending on whether you want to perform the operation at the group level or pair level:

**Group level:**

Select a group, click , and select **Split**.

**Pair level:**

- a. Select a group and click  to open its **Details** view.
  - b. Click on the number next to **Clone Pairs**.
  - c.
    - Select one or more pairs, click , and select **Split**.
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. Optional: To attach session options to the operation, click **Advanced Options**, 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Terminating clone copy sessions

### Before you begin


- TimeFinder/Clone requires Engenuity version 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.

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.



### Procedure

1. To split clone volume pairs:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Device Groups**.
  3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
  4. Do one of the following, depending on whether you want to perform the operation at the group level or pair level:

**Group level:**


Select a group, click , and select **Terminate**.

**Pair level:**

- a. Select a group and click  to open its **Details** view.
  - b. Click on the number next to **Clone Pairs**.
  - c.
    - Select one or more pairs, click , and select **Terminate**.
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 **Advanced Options**, and select any number of [Clone copy session options](#) on page 365.
  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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing clone pairs


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
4. Select a group and click  to open its **Details** view.
5. Click on the number next to **Clone Pairs**.

The following properties display:

- **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.

The following controls are available:

-  —[Viewing clone pair details](#) on page 365
- **Create Pairs**—[Creating clone copy sessions](#) on page 355
- **Activate**—[Activating clone copy sessions](#) on page 357
- **Recreate**—[Recreating clone copy sessions](#) on page 358
- **Split**—[Splitting clone volume pairs](#) on page 362
- **Restore**—[Restoring data from target volumes](#) on page 361
- **Set Mode**—[Modifying clone copy sessions](#) on page 360
- **Create Snapshot**—[Creating clone snapshots](#) on page 359
- **Terminate**—[Terminating clone copy sessions](#) on page 363

## Viewing clone pair details

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Clone** tab to open the **TimeFinder Clone** list view.
4. Select a group, and click ⓘ to open its **Details** view .
5. Click on the number next to **Clone Pairs**.

6. Select a pair and click ⓘ to open its **Details** view.

The following properties display:

- **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** — (this property is displayed by clicking ⓘ). 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. (this property is displayed by clicking ⓘ).
- **Timestamp** — Date and time the pair was created. (this property is displayed by clicking ⓘ).

## 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 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

**Table 5** TimeFinder/Clone session options (continued)

Session option	Description	Available with action
		Recreate
Restored	With the verify command, verifies that the copy sessions are in the Restored state. With the terminate command, terminates a restored VP Snap session.	Verify Terminate
Star	Targets the action at volumes in SRDF/Star mode.	Create Recreate Establish Activate Restore Split Terminate
Symforce	Forces an operation on the volume pair including pairs that would be rejected. Use caution when checking this option because improper use may result in data loss.	Terminate

## 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

### Before you begin

TimeFinder/Snap requires Enginuity version 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.

The **TimeFinder/Snap** dashboard provides you with a single place to monitor and manage TimeFinder/Snap sessions on a storage system.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.

The following properties display:

- **Device 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. Property values: Regular, R1, R2, or R21.
- **Group Valid**—Whether the group is valid or invalid.

Click ⓘ and click the number next to **Snap Pairs** to view the associated snap pairs (see [Viewing snap pairs](#) on page 376).

Click ⓘ to click the number next **Storage Groups** to view the associated storage groups.

The following controls are available:

- **Create Pairs**—[Creating virtual copy sessions](#) on page 369
- **Activate**—[Activating virtual copy sessions](#) on page 370
- **Terminate**—[Terminating virtual copy sessions](#) on page 375
- **Restore**—[Restoring virtual copy sessions](#) on page 374
- **Recreate**—[Recreating virtual copy sessions](#) on page 373



- **Establish**—[Creating snapshots](#) on page 371
- **Duplicate**—[Duplicating virtual copy sessions](#) on page 372

## 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

---

### Note

- TimeFinder/Snap requires Enginuity version 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. To do this, 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 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:


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level.

Group level:

- a. Select a group, and click **Create Pairs**.
- b. Select a source type and a target type.

Pair level:

- a. Select a group, and click  to open its **Details** view.
  - b. Click on the number next to **SnapPairs**.
  - c. Select one or more pairs and click **Create Pairs**.
  - d. Click **Set Pairs** to open the **Set Pairs** dialog box.
  - e. Select a source volume and a target volume, and click **Add** to make them a pair. Repeat this step as required.
  - f. Click **OK** to return to the **Create Sessions** dialog box.
5. Click **Advanced Options** to set the advanced options as described next.
    - Select a **Pairing Type**. If you are not using the **Pairing Type** option, leave this field set to **None**.
    - To attach Session Options to the operation, select any number of [options](#).
  6. Do one of the following:
    - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
    - 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.

---

### Note

TimeFinder/Snap requires Enginuity version 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:


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

- Select a group, and click **Activate**.
- Select a source type and a target type.

Pair level:

- a. Select a group, and click  to open its **Details** view.
- b. Click on the number next to **Clone Pairs**.

- c. Select one or more pairs, and click **Activate**.
5. Click **Advanced Options** to set the advanced options. To attach session options to the operation, select any number of [options](#).
6. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Creating snapshots

### Before you begin

TimeFinder/Snap requires Engineuity OS 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.


This procedure explains how to create and immediately activate virtual copy sessions.

To create a snapshot:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

#### Group level:


Select a group, click , and select **Create Snapshot**.

Select a source type and a target type.

#### Pair level:

- a. Select a group, and click  to open its **Details** view.

b.

Select one or more pairs, click , and select **Create Snapshot**.

5. Click **Advanced Options** to set the advanced options.

Setting Advanced options:

- a. Select one of the following for **Pairing Type**. 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.
- b. To attach Session Options to the operation, select **Advanced Options** and select any number of [options](#).

c. Click **OK**.

6. Do one of the following:

- Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## 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 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 session:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

Group level:

Select a group, click , and select **Duplicate**.

Select a source type and a target type.

Pair level:

a. Select a group, and click  to open its **Details** view.

b.

Select one or more pairs, click , and select **Duplicate**.

5. To attach Session Options to the operation, select **Advanced Options** and select any number of [options](#).

6. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Recreating virtual copy sessions

### Before you begin


- TimeFinder/Snap requires Enginuity version 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 5876 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.

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.

### Procedure

1. To recreate virtual copy sessions:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Device Groups**.
  3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.
  4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

#### Group level:

Select a group, click , and select **Recreate**.

Select a source type and a target type.

#### Pair level:

- a. Select a group, and click  to open its **Details** view.

b.

Select one or more pairs, click , select **Recreate** to open the **Recreate** dialog box.

5. To attach Session Options to the operation, select **Advanced Options** and select any number of [options](#).
6. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Restoring virtual copy sessions

### Before you begin

- TimeFinder/Snap requires Enginuity version 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.
- With Enginuity 5876 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).


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.


### Procedure


1. To restore virtual copy sessions:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Device Groups**.
  3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.
  4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

#### Group level:

Select a group, click , and select **Restore**.  
Select a source type and a target type.

#### Pair level:

- a. Select a group, and click  to open its **Details** view.
- b.

Select one or more pairs, click , and select **Restore** to open the **Restore** dialog.

5. Select the **Restore Type**.

Restore operations can be used to copy target data to another device (full restore), or back to the original source device (incremental restore). In the case of a full restore, the original session terminates and a copy session to the target of the restore starts. In the case of an incremental restore, the original session

copy direction is reversed and changed data is copied from the target device to the source device. Restore operations require that the original session is differential and the source device is fully copied.

6. If performing a **Full** restore, click **Set Pairs** to open the **Set TimeFinder Snap Pairs** dialog from which you can select the volumes to use in the operation.
7. To attach Session Options to the operation, select **Advanced Options** and select any number of [options](#).
8. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Terminating virtual copy sessions

### Before you begin

TimeFinder/Snap requires Enginuity version 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Snap operations are mapped to their TimeFinder/SnapVX equivalents.

This procedure explains how to terminate an active virtual copy session at any time.

### Procedure


1. To terminate virtual copy sessions:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Device Groups**.
  3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.
  4. Do one of the following, depending on whether you want to perform the operation at the group level or pair level:
 

**Group level:**

Select a group and select **Terminate**.



Select a source type and a target type.

**Pair level:**

    - a. Select a group, and click  to open its **Details** view.
    - b. Select one or more pairs and select **Terminate**.
5. To attach Session Options to the operation, select **Advanced Options** and select any number of [options](#).
6. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing snap pair details

### Procedure


1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.
4. Select a group, and click  to open its **Details** view .
5. Click on the number next to **Snap Pairs**.
6. Select a pair and click  to open its **Details** view.

The following properties display:

- **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.
- **State**—Session state of the pair.
- **Snap Pool** — The name of the snap pool.
- **Percent Copied** — The percentage of copying that is complete.
- **Timestamp** — Date and time the snapshot was created.

## Viewing snap pairs


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Snap** tab to open the **TimeFinder Snap** list view.
4. Select a group and click  to open its **Details** view.
5. Click on the number next to **Snap Pairs**.

The following properties display:

- **Source Volume** — The name of the source volume.
- **Target Volume** — The name of the target volume.
- **Source LDev** — The logical name of the source volume.
- **Target LDev** — The logical name of the target volume.
- **State** — The session state of the pair.

The following controls are available:

-  — [Viewing snap pair details](#) on page 376
- **Create Pairs**—[Creating virtual copy sessions](#) on page 369
- **Activate**—[Activating virtual copy sessions](#) on page 370



- **Terminate**—[Terminating virtual copy sessions](#) on page 375
- **Detach**—[Viewing clone pairs](#) on page 364
- **Attach**—[Viewing clone pairs](#) on page 364
- **Duplicate**—[Duplicating virtual copy sessions](#) on page 372
- **Create Snapshot**—[Creating snapshots](#) on page 371
- **Recreate**—[Recreating virtual copy sessions](#) on page 373
- **Restore**—[Restoring virtual copy sessions](#) on page 374

## 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

## Set TimeFinder Snap Pairs dialog box

When creating, activating, restoring, or establishing a TimeFinder/Snap pairs, this dialog box allows you to define the pairs used in the operation.

### Procedure

1. To define the pairs:
  1. Select the **Source Volumes** and **Target Volumes** and click **Add** to move them to the **Selected Pairs** table.
  2. Click **OK**.

## Managing TimeFinder/Mirror sessions

### Before you begin

- TimeFinder/Mirror requires Enginuity version 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.

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

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Mirror** tab to open the **TimeFinder Mirror** list view.

The following properties display:

- **Device Group**—Groups containing volumes using TimeFinder/Mirror.
- **Standard**—The number of standard volumes in the group.
- **BCVs**— 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. Property values are: RDF1, RDF2, RDF21, and Regular
- **Group Valid** —Indicates whether the group is valid. Property values are: Yes and No.

Click  and click the number next to **Mirror Pairs** to view the associated mirror pairs (see [Viewing snap pairs](#) on page 376).

Click  to click the number next **Storage Groups** to view the associated storage groups.

The following controls are available:

- **Create Snapshot**—[Creating Snapshots](#) on page 379
- **Restore**—[Restoring BCV pairs](#) on page 380
- **Split**—[Splitting BCV pairs](#) on page 381

- **Cancel**—[Cancelling BCV pairs](#) on page 381

## Creating Snapshots

### Before you begin

- TimeFinder/Mirror requires Engenuity version 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.


### Procedure

1. To create snapshots:
1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Mirror** tab to open the **TimeFinder Mirror** list view.
4. Do the following, depending on whether you want to perform the operation at the group level or pair level:

#### Group level:

Select a device group, and click **Create Snapshot** to

#### Pair level:

- a. Select a device group, and click  to open its **Mirror Sessions List** view.
  - b. Select one or more pairs, click **Create Snapshot** to open the **Create Snapshot - Mirror Pair** dialog.
5. 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.
  6. If performing a full establish at the pair level, do the following:
    - a. Click **Set Pairs** to open the **Set TimeFinder Mirror Pairs** dialog.
    - 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.
  7. To attach session options to the operation, select **Advanced Options** and select any number of [options](#).
  8. Do one of the following:
    - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Restoring BCV pairs

### Before you begin

TimeFinder/Mirror requires Enginuity version 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Mirror operations are mapped to their TimeFinder/SnapVX equivalents.

This procedure explains how to copy data from the BCV volumes to the standard volumes.


### Procedure

1. To restore BCV pairs:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Device Groups**.
  3. Click the **TimeFinder Mirror** tab to open the **TimeFinder Mirror** list view.
  4. 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**.

**Pair level:**

    - a. Select a device group, and click  to open the **Mirror Sessions List** view.
    - b. Select one or more pairs, and click **Restore**.
5. Select a **Restore Type**:
  - **Incremental**—Copies to the standard volume only the new data that was updated on the BCV volume while the BCV pair was split.
  - **Full**—Copies the entire contents of the BCV volume to the standard volume.
6. If performing a full establish at the pair level, do the following:
  - a. Click **Set Pairs** to open the **Set TimeFinder Mirror Pairs** dialog.
  - 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.
7. To attach session options to the operation, select **Advanced Options** and select any number of [options](#).
8. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Splitting BCV pairs

### Before you begin

TimeFinder/Mirror requires Enginuity version 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Mirror operations are mapped to their TimeFinder/SnapVX equivalents.

This procedure explains how to split paired volumes to where each holds separate valid copies of the data.


### Procedure

1. To split BCV pairs:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Device Groups**.
  3. Click the **TimeFinder Mirror** tab to open the **TimeFinder Mirror** list view.
  4. 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 **Split**.

#### Pair level:

- a. Select a device group, and click  to open the **Mirror Sessions List** view.
  - b. Select one or more pairs, and click **Split**.
5. To attach session options to the operation, select **Advanced Options** and select any number of [options](#).
  6. Do one of the following:
    - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Cancelling BCV pairs

TimeFinder/Mirror requires Enginuity version 5876. On storage systems running HYPERMAX OS 5977 or higher, TimeFinder/Mirror operations are mapped to their TimeFinder/SnapVX equivalents.

To cancel the relationship between volumes in a BCV pair:


### Procedure

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 **Cancel**.

Pair level:

- a. Select a device group and click  to open the **Mirror Sessions List** view.
  - b. Select one or more pairs and click **Cancel**.
4. To attach session options to the operation, select **Advanced Options** and select any number of [options](#).
  5. Do one of the following:
    - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing mirror pairs


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Device Groups**.
3. Click the **TimeFinder Mirror** tab to open the **TimeFinder Mirror** list view.
4. Select a group and click  to open its **Details** view.
5. Click on the number next to **Mirror Pairs**.

The following properties display:

- **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.
- **Pair State**—The session state of the pair.
- **Timestamp**—Date and time the snapshot was created.




The following controls are available:

-  —[Viewing mirror pair details](#) on page 382
- **Create Snapshot**—[Creating Snapshots](#) on page 379
- **Restore**—[Restoring BCV pairs](#) on page 380
- **Split**—[Splitting BCV pairs](#) on page 381
- **Cancel**—[Cancelling BCV pairs](#) on page 381

## Viewing mirror pair details

### Procedure

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  to open its **Mirror Pairs List** view.
4. Select a pair and click  to open its Details view.
5. Click on the number next to **Mirror Pairs**.
6. Select a pair and click  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.
  - **State**—Session state of the pair.
  - **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

**Table 7** TimeFinder/Mirror session options (continued)

<b>Session option</b>	<b>Description</b>	<b>Available with action</b>
Differential	Used with either the Copy or Precopy option to create an SDDF session for maintaining changed track information. This must be used when creating copy sessions on which you plan on issuing a Restore action.	Create Establish
Force	Overrides any restrictions and forces the operation, even though one or more paired volumes may not be in the expected state. Use caution when checking this option because improper use may result in data loss.	Create Establish Activate Restore Split Terminate
Not Ready	Sets the target volumes as Not Ready. Upon completion of a split action, the target volumes are set as Not Ready. When a restore is initiated, the standard volumes are set as Not Ready.	Split Full Restore Incremental Restore
Optimize	Optimizes volume pairings across the local 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	Split Full Restore Incremental Restore



**Table 7** TimeFinder/Mirror session options (continued)

Session option	Description	Available with action
	defaults to not propagate the data to the remote mirror of the RDF volume.	
Reverse	With a split operation, initiates a reverse data copy from one or more fixed BCV mirrors to the first (moving) mirror of the BCV upon the completion of the split operation. With an establish or restore operation, requests a verification check that the BCV's fixed mirror has valid data. If at establish or restore time you anticipate a need to perform future BCV reverse split operations, you must apply a reverse establish or restore so that no invalid tracks on the fixed BCV mirror become used.	Full Establish Incremental Establish Split Full Restore Incremental Restore
Star	Targets the action at volumes in SRDF/Star mode.	Full Establish Restore Split Cancel
SymForce	Forces an operation on the volume pair including pairs that would be rejected. Use caution when checking this option because improper use may result in data loss.	Full Establish Incremental Establish Split Full Restore Incremental Restore

## Setting TimeFinder/Mirror pairs

When establishing or restoring TimeFinder/Mirror pairs, this dialog box allows you to define the pairs used in the operation.

### Procedure

1. To define the pairs:
  1. Select the **Source Volumes** and **Target Volumes** and click **Add** to move them to the **Selected Pairs** table.
  2. Click **OK**.

## 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 view provides a single place from you can manage TimeFinder SnapVX snapshots and their associated storage groups.

**secure snaps**—These are Snap VX snapshots that can't be deleted before the expiry time set by the StorageAdmin. Users can create a Secure snapshot or set Secure status on an existing snapshot. Once the retention time has expired, the Secure snapshot will be automatically terminated unless there is a linked device or an active restore session is ongoing. The expiry time on a Secure snapshot can be changed but the time can only be moved forward from the expiry time originally set. This feature requires an array running the HYPERMAX OS 5977 Q1 2017 Service Release or higher.

---

#### Note

Secure snapshots may only be terminated after they expire or by customer-authorized support. Please refer to Knowledge Base article 498316 for additional information.

---

**Time To Live** — From Unisphere 8.4 onwards, users can now specify a SnapVX snapshot's time to live in hours as well as days. Previously only days could be specified.

#### 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 view:


1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.

#### TimeFinder/SnapVX view

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.
- **Last Creation Time**—Date/time the most recent snapshot was created.

The following controls are available:

-  —displays a properties panel listing the following properties: Storage Group, Capacity (GB), Number of Snapshots and SRP.
- **Create**—[Creating snapshots](#) on page 387
- **Modify**—[Modifying TimeFinder SnapVX snapshots](#) on page 389
- **Restore**—[Restoring snapshots](#) on page 393
- **Link**—[Linking to snapshots](#) on page 390

- **Unlink**—[Unlinking from snapshots](#) on page 392
- **Relink**—[Relinking to snapshots](#) on page 391
- **Set Mode**—[Setting copy mode for snapshots](#) on page 396
- **Set Time to Live** —[Setting snapshots to automatically terminate](#) on page 393
- **Set Secure**—[Setting "Secure" status on an existing snapshot](#) on page 394
- **Terminate**—[Terminating snapshots](#) on page 395

## Creating 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.
- The Secure snapshot feature requires the HYPERMAX OS 5977 Q1 2017 Service Release or higher.
- You can perform this operation from the following : **TimeFinder/SnapVX** view , **Storage** view, or **Data Protection** dashboard. Depending on the location from which you are performing this operation, some of the following steps may not apply.

This procedure explains how to create TimeFinder SnapVX snapshots.

---

### Note

Secure snapshots may only be terminated after they expire or by customer-authorized support. Please refer to Knowledge Base article 498316 for additional information.

---

To create snapshots:

### Procedure

1. Select the storage system.
2. Do the following, depending on the location from which you want to perform the procedure:

#### TimeFinder/SnapVX view:

- a. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
- b. Select a storage group and click **Create** to open the **Create Snapshot** dialog.

#### Storage Groups view:

- a. Select **STORAGE > Storage Groups** to open the **Storage Groups** view.
- b. Select the storage group and click **Protect** to open the **Protect Storage Group** wizard.
- c. If not already selected, select **Point In Time Protection Using SnapVX**.
- d. Click **Next**.

#### Data Protection dashboard:

- a. Select **Replication** to open the **Data Protection** dashboard.
  - b. Click **CREATE SNAPSHOT**.
3. Select whether to create a new snapshot or reuse an existing snapshot.
  4. If reusing an existing snapshot, select it from the list. When using this method, 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.**

---

5. Choose an expiry type from the drop-down menu. The options are:
  - **None** — If no automatic expiry time is set the snapshot will need to be manually deleted.
  - **Time to live** — Once the time you set has expired, the snapshot will be automatically terminated, provided that it is not linked to any target volumes. If an expired snapshot is linked, the system waits until the last link has been removed before terminating the snapshot. To override this behavior and terminate the snapshot, select the **Force** option under the **Advanced Options** link.
6. If you chose **Time to live** as the protection type, use the **Days** and **Hours** drop-down menus to set the snapshot's expiry time.
7. Click **Advanced Options** to see the advanced options. They are:
  - **Enable Secure Snaps** — Select this option to set a Secure snapshot that can't be deleted before the expiry time you set. Once you tick the **Secure** checkbox the **Days** and **Hours** drop-down menus will appear and you can use these to set the snapshot's expiry time. Once the retention time has expired, the Secure snapshot will be automatically terminated unless there is a linked device or an active restore session is ongoing. StorageAdmins can choose to move the retention time forward.
  - **Both Sides** — Select this option to create a snapshot at both sides of an SRDF pairing simultaneously. The following limitations apply:
    - A consistent snapshot on both sides is only allowed when the SRDF pairs exist on the source Storage Group volumes in Synchronous RDF mode and the SRDF pair state is Synchronous.
    - A consistent snapshot on both sides is only allowed when the SRDF pairs are in Active SRDF mode and the SRDF pair state is ActiveActive or ActiveBias.
    - A mixture of R1 and R2 devices is not allowed.
    - All the RDF devices in the SG must be in same RDF group.
    - Concurrent RDF is not supported.

- For cascaded SRDF setups, the Both Sides option is supported by selected Storage Group and the next immediate hop, but not the subsequent hops.
  - **Enable Force Flag** — Select this option to force the operation even though one or more volumes may not be in the normal, expected states.
8. Click **Next**.
  9. Choose 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. This option can be used to create a recurring daily SnapVX snapshot for a given time. In the event of a failed recurring snapshot, an Alert will be raised to notify the user. The schedule continues to run in the event of a failed snapshot, issuing alerts to the user. The alerts list view will retain a record of the failed snapshots (unless the alert is deleted). A warning level alert will be issued. There will not be an end date for the schedule specified when setting it up, so you will need to cancel the schedule manually, if desired. For more information, refer to [Scheduling jobs](#) on page 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List** and click **Run Now** to perform the operation now.

## Modifying TimeFinder SnapVX snapshots

### Before you begin

- To perform this operation, you must be a StorageAdmin.
- The Secure snapshot feature requires the HYPERMAX OS 5977 Q1 2017 Service Release or higher.

To modify TimeFinder SnapVX snapshots:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3. Select a snapshot and click **Modify** to open the **Edit Snapshot** dialog.
4. Enter the new name for the snapshot.
5. Choose 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. This option can be used to create a recurring daily SnapVX snapshot for a given time. In the event of a failed recurring snapshot, an Alert will be raised to notify the user. The schedule continues to run in the event of a failed snapshot, issuing alerts to the user. The alerts list view will retain a record of the failed snapshots (unless the alert is deleted). A warning level alert will be issued. There will not be an end date for the schedule specified when setting it up, so you will need to cancel the schedule manually, if desired. For more information, refer to [Scheduling jobs](#) on page 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List** and click **Run Now** to perform the operation now.

## Linking to 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 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.
- The SnapVX link storage group dialog is updated to always create CKD devices when the New storage group target name radio button is selected.

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.

Linking a Storage Groups snapshot after the SG volumes have been subsequently expanded will pick volumes to link to by using the volume size at the time of the snapshot being taken.

### Procedure

1. To link to snapshots:
  1. Select the storage system.
  2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
  3. Select the storage group and click **Link**.
  4. Select the **Snapshot Name**.
  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. Optional: Modify the default name for the new storage group.
  7. Click **Advanced Options** to continue setting the advanced options, as described next.

Setting Advanced options:

To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.

Setting advanced options

- To create a permanent, full-time copy of the data on the target volume's SRP, select **Copy**. Selecting **Copy** enables the **Remote** option.
  - To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.
  - To specify that the operation is for devices in STAR mode, select **Star**.
  - Optional: Uncheck the **Compression** check box to turn off Compression. Compression is only allowed on All Flash systems running the HYPERMAX OS 5977 Q3 2016 Service Release or higher.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Relinking to snapshots



### 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.

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.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3.  Select the storage group, click , and select **Relink** to open the **Relink** dialog box.
4. Select the link target Storage group and the Snapshot Name.
5. Click **Advanced Options** to continue setting the advanced options, as described next.

Setting Advanced options:

- To create a permanent, full-time copy of the data on the target volume's SRP, select **Copy**. Selecting **Copy** enables the **Remote** option.
  - To specify that the operation is for devices in STAR mode, select **Star**.
  - To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Unlinking from 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.


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 page 391.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3.
 



 Select the storage group, click , and select **Unlink** to open the **Unlink** dialog box.
4. Select the **Snapshot Name**.
5. Click **Advanced Options** to continue setting the advanced options, as described next.
 

Setting Advanced options:

To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.

To specify that the operation is for devices in STAR mode, select **Star**.

To force the operation when the operation would normally be rejected, select **SymForce**.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.



## Restoring 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 view. You can also perform this operation from other locations in the interface. Depending on the location, some of the steps may not apply.

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.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3. Select the storage group and click **Restore**.
4. Select the **Snapshot Name** and **Creation Date** (0 is the latest).
5. Click **Advanced Options** to continue setting the advanced options, as described next.

Setting Advanced options:

To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.

To specify that the operation is for devices in STAR mode, select **Star**.

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting snapshots to automatically terminate


### 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.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.

3.

Select the storage group, click , and select **Set Time to Live** to open the **Set Time to Live** dialog box.

4. Select the **Snapshot Name** and **Creation Date**.

5. Select the amount of days and hours you want the snapshot to exist for.

Once the time has expired, the snapshot is automatically terminated, provided that it is not linked to any target volumes. If an expired snapshot is linked, the system will wait until the last link has been removed before terminating the snapshot. To override this behavior, select the **Force** option, which will allow the system to terminate the snapshot regardless of whether it is linked. To remove the Time to Live attribute, select **None**.

6. Click **Advanced Options** to continue setting the advanced options, as described next.

Setting Advanced options:

To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.

To specify that the operation is for devices in STAR mode, select **Star**.

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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting "Secure" status on an existing snapshot

### Before you begin

To perform this operation, you must be a StorageAdmin.

The Secure snapshot feature requires the HYPERMAX OS 5977 Q1 2017 Service Release or higher.

This procedure explains how to set "Secure" status on an existing snapshot. It can also be performed by clicking on a storage group in the **TimeFinder SnapVX** view to open the **Snapshots** view.

---

### Note

Secure snapshots may only be terminated after they expire or by customer-authorized support. Please refer to Knowledge Base article 498316 for additional information.


---

### Procedure

1. Select the storage system.

2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.

3.

Select the storage group, click , and select **Terminate** to open the **Terminate** dialog box.

4. Select the name of an existing snapshot and then use the **Days** and **Hours** drop-down menus to set the expiry time.

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.


## Terminating 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 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.
- If the snapshot is Restored, then this action terminates the restore session. If you want to terminate the snapshot, the dialog and action have to be executed again.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3.
 



 Select the storage group, click , and select **Terminate** to open the **Terminate** dialog box.
4. Select the **Snapshot Name**.
5. Click **Advanced Options** to continue setting the advanced options, as described next.

Setting Advanced options:

To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.

To specify that the operation is for devices in STAR mode, select **Star**.

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.**

---

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.


## Setting copy mode for 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.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION** > **Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
- 3.

Select the storage group, click , and select **Set Mode** to open the **Set Mode** dialog box.

4. Select the **Snapshot Name**.
5. Select a new 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.
  - **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.
6. Click **Advanced Options** to continue setting the advanced options, as described next.

Setting Advanced options:

To force the operation even though one or more volumes may not be in the normal, expected state(s), select **Force**.

To specify that the operation is for devices in STAR mode, select **Star**.

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing snapshots

### Before you begin


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

This procedure explains how to view and manage snapshots of a storage group.

### Procedure


1. Select the storage system.

2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.

3. Select a storage group, click  and click on the number next to **Number of Snapshots**

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**—Indication whether the snapshot is linked to another storage group. A checkmark indicates that the snapshot is linked.
- **Restored**—Indication whether the snapshot is restored to the source. A checkmark indicates that the snapshot is restored.
- **Time To Live**—time the snapshot has to live.
- **Secured**—Whether the snapshot is Secured or not. A checkmark indicates that the snapshot is Secured, a dash indicates that it isn't. "Expired" indicates that the snapshot was Secured but is now expired.
-  —[Viewing snapshot details](#) on page 397
- **Create**—[Creating snapshots](#) on page 387
- **Modify**—[Modifying TimeFinder SnapVX snapshots](#) on page 389
- **Restore**—[Restoring snapshots](#) on page 393
- **Link**—[Linking to snapshots](#) on page 390
- **Unlink**—[Unlinking from snapshots](#) on page 392
- **Relink**—[Relinking to snapshots](#) on page 391
- **Set Mode**—[Setting copy mode for snapshots](#) on page 396
- **Set Time to Live**—[Setting snapshots to automatically terminate](#) on page 393
- **Set Secure**—[Setting "Secure" status on an existing snapshot](#) on page 394
- **Terminate**—[Terminating snapshots](#) on page 395

## Viewing snapshot details

### Before you begin

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

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3. Select a storage group, click  and click on the number next to **Number of Snapshots**
4. Select a snapshot and select  to open the snapshot **Details** view.

The Snapshot **Details** view allows you to view and manage a snapshot.

#### Properties panel

The following properties display:

- **Name**—Name of the snapshot.
- **Storage Group Name**—Name of the snapshot.
- **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. assigns generation numbers to the snapshots in the order in which they were created (latest = generation 0, previous incrementing by one).
- **Creation Time**—Date and time the snapshot was created.
- **Expiry Date**—Date and time the snapshot is set to automatically terminate if either "Secure" or "Time to Live" has been set. If the snapshot is not set to automatically terminate, this field displays N/A.
- **State**—Snapshot state.
- **Secured**—Indicates whether the snapshot is Secured or not. A checkmark indicates that the snapshot is Secured, a dash indicates that it isn't. "Expired" indicates that the snapshot was Secured but is now expired.



There are also links to views displaying objects (Source Volumes, Links and SRP) 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** opens a view listing the links associated with the snapshot.

## Viewing snapshot links

### Before you begin

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

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3. Select a storage group, click  and click on the number next to **Number of Snapshots**
4. Select a snapshot and select  to open the snapshot **Details** view.
5. Click on the number next to **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.
- **Snapshot Timestamp**—Date and time the snapshot was created.
- **Link Timestamp**—Date and time the link was created.

The following controls are available:

- ⓘ —Displays a properties panel listing the following properties: Source Storage Group and Linked Volumes.
- **Unlink**—[Unlinking from snapshots](#) on page 392
- **Relink**—[Relinking to snapshots](#) on page 391

## Viewing snapshot link details

### Before you begin

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

### Procedure

1. To view snapshot link details:
1. Select the storage system.
2. Select **DATA PROTECTION** > **Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3. Select a storage group, click ⓘ and click on the number next to **Number of Snapshots**
4. Select a snapshot and select ⓘ to open the snapshot **Details** view.
5. Click on the number next to **Links** to open the snapshot **Links** list view.
6. Select a snapshot and select ⓘ to open the snapshot links **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.
- **Linked Volumes**—Name of the linked volume(s).
- **State**—Snapshot state.
- **Flags (FCMD)**—Snapshot flags. Possible values are: Failed, Copied, Modified, Defined (FCMD).

## Viewing snapshot source volumes

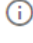
### Before you begin

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

This view displays SnapVX ICDP snapshots created from the Mainframe product. Management of these snapshots is not supported.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION** > **Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3. Select a storage group, click ⓘ and click on the number next to **Number of Snapshots**


4. Select a snapshot and select  to open the snapshot **Details** view.
5. Click on the number next to **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 are displayed:

- **Name**—Name of volume
- **State**—Snapshot state.
- **Creation Date**—Date and time the snapshot was created.
- **Failed**—Indication of failure.
- **Linked**—Indication of link status.
- **Restored**—Indication of restoration status.

The following controls are available:




-  —[Viewing snapshot source volume details](#) on page 400
- **Restore**—[Restoring snapshots](#) on page 393
- **Link**—[Linking to snapshots](#) on page 390
- **Relink**—[Relinking to snapshots](#) on page 391
- **Unlink**—[Unlinking from snapshots](#) on page 392
- **Set Mode**—[Setting copy mode for snapshots](#) on page 396
- **Set Time to Live**—[Setting snapshots to automatically terminate](#) on page 393
- **Set Secure**—[Setting "Secure" status on an existing snapshot](#) on page 394
- **Terminate**—[Terminating snapshots](#) on page 395

## Viewing snapshot source volume details

### Before you begin

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

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION** > **Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3. Select a storage group, click  and click on the number next to **Number of Snapshots**
4. Select a snapshot and select  to open the snapshot **Details** view.
5. Click on the number next to **Source Volumes** to open the snapshot **Source Volumes** list view.
6. Select the volume and click  to open the snapshot source volume **Details** view.



The snapshot source volume **Details** view allows you to view and manage the source volume in a snapshot.

The following properties display:




- **Name**—Name of the volume.
- **State**—Snapshot state.
- **Secured**—Snapshot secured indication.
- **Flags**—Snapshot flags. Possible values are: Failed, Link, Restore, GCM, Type (FLRGFT).
- **Capacity (GB)**—Capacity of the volume.
- **Tracks**—Number of source tracks that the host has not yet overwritten.
- **Track Size**—Track size in bytes.
- **Linked Volumes**—Linked volumes.

## Viewing snapshot source volume linked volumes

### Before you begin

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

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION** > **Storage Groups** and click on the **SnapVX** tab to open the **TimeFinder/SnapVX** view.
3. Select a storage group, click  and click on the number next to **Number of Snapshots**
4. Select a snapshot and select  to open the snapshot **Details** view.
5. Click on the number next to **Source Volumes** to open the snapshot **Source Volumes** list view.
6. Select the volume and click  to open the snapshot source volume **Details** view.
7. Click the number next to **Linked 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.
- **Snapshot Timestamp**—Date and time the snapshot was created.
- **LinkTimestamp**—Date and time the link was created.

The following controls are available:

- **Unlink**—[Unlinking from snapshots](#) on page 392

- **Relink**—[Relinking to snapshots](#) on page 391

## RBAC roles for performing local and remote replication actions

The table below details the roles needed to perform TimeFinder SnapVX local and remote replication actions.

### Note

Unisphere for PowerMax does not support RBAC Device Group management.

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

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

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

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

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

## Managing remote replication sessions

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

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

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

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

## Creating SRDF connections

This task provides a mechanism to make a connection to storage array that is currently not visible to the Unisphere server and to bring the connected array into Unisphere as remote.


Before you begin:

The physical connectivity and zoning must be in place before undertaking this task.

To create SRDF connections:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF Groups**.
- 3.

Select an SRDF group, click , and select **Create SRDF Connection** to open the **Create SRDF Connection** wizard.

4. On the **Local** page, specify the following information:
  - Type a value for the SRDF group label.
  - Select a SRDF Group Number from the list of unused RDFG numbers for the local array.
  - From the list, select a local port to be used by the new SRDF Group.
5. (Optional) Click **NEXT**.
6. On the **Remote** page, specify the following information:
  - Select **Scan** to scan the SRDF SAN for the port selected on the local page.
  - Select an Array ID from the list.
  - Type a value for the SRDF Group Number. This is not selectable as there is no knowledge of the remote candidate array's used RDFG numbers at this point.
  - From the list, select a remote port to be used by the new SRDF Group.
7. (Optional) Click **NEXT**.
8. On the **Summary** 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.

9. Do one of the following:

- Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

### Results

A SRDF group has been created with a single port on each side. After creation, further SRDF group changes can be performed using Unisphere functionality.

## Creating SRDF pairs

### Before you begin

Creation of an SRDF pair can be blocked 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*.

You can 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 OS 5876 or higher, and the standard volume must be on a storage system running Enginuity OS 5876.

Meta volumes are supported on storage systems running Enginuity OS 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:

Unisphere provides support for creating RDF pairs in a concurrent RDF in a SRDF/Metro configuration resulting in one Metro RDF mirror and one Async or Adaptive Copy RDF mirror.

**Note**

The following restrictions apply:

- Adding a Metro RDF mirror when the device is already part of an SRDF/Metro configuration.
- Adding a Metro RDF mirror when the device is already an R2 device.
- Adding a non-Metro RDF R2 mirror to a device that has a Metro RDF mirror.
- Adding a Metro RDF mirror when the non-Metro RDF mirror is in Synchronous mode.
- Adding a non-Metro RDF mirror in Synchronous mode when the device is already part of an SRDF/Metro configuration
- Operations that make the Metro RDF mirror RW on the RDF link are not allowed if the Metro device is the target of the data copy from the non-Metro RDF mirror.
- Operations that make the non-Metro RDF mirror RW on the RDF link and result in the data copy to the Metro device are not allowed if the Metro RDF mirror is RW on the RDF link.
- The Create Pair - Invalidate R1 operation is not allowed on the non-Metro RDF mirror if it results in a Metro device becoming write-disabled (WD).

R1	R21	R2
5977	5977	5977
5977	5977	5876
5977	5876	5876
5977	5876	5977
5876	5876	5977
5876	5977	5977
5876	5977	5876

If the RDF interaction includes a storage system running HYPERMAX OS 5977 or higher, then the other storage system must be running Enginuity OS 5876 or higher.

It is possible to create a SRDF/Metro device pair when SRDF/Metro exist in a current group or an empty SRDF group exists on the storage device.

CKD devices are not supported by SRDF/Metro.

Only CKD storage groups are selectable if the volumes chosen are of that emulation.

If Local or Remote storage system is running Enginuity OS 5876, only Bound TDEVs are supported, and this requires the selection of a thin pool.

Adding to Storage Groups will list SGs which are either empty or not a parent (i.e. child or standalone). SGs which already contain devices must have those devices in the SRDF group which the wizard is being run against, and have the devices of the same SRDF polarity (R1s or R2s).

This procedure supports adding SRDF pairs to a SRDF/Metro group.

To create an SRDF pair:

## Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > 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 Mirror Type to apply to the local devices.
5. Select SRDF Mode.
6. Select Adaptive Copy Mode option for (Disk / Write Pending) (storage systems running Enginuity OS 5876 only)
7. Select one of the following options:
  - **Invalidate R1**- Invalidates the source R1 device(s) so that a full copy can be initiated from the remote mirror.
  - **Invalidate R2**- Invalidates the target R2 device(s) so that a full copy can be initiated from the remote mirror.
  - **Establish**- Begins a full copy from the source to the target, synchronizing the dynamic SRDF pairs in the device file.
  - **Restore**- Begins a full copy from the target to the source, synchronizing the dynamic SRDF pairs in the device file.
  - **Format**- No data resynchronization is done between source and target dynamic SRDF pairs in the device file after all tracks are cleared on what will become the R1 and R2 side.
8. Optional: Select **No WD** - Bypasses the check that ensures that the target of the operation is not writable by the host.
9. Click **NEXT** to go to the **Local Volumes** page.
10. If you wish to do manual selection for local devices, turn Automatic Selection off.
11. Select the thin pool name.
12. Specify criteria to find the volumes of interest, and choose volumes.
13. Click **Add to Storage Group** checkbox and select a storage group.
14. Click **NEXT** to go to the **Remote Volumes** page.
15. If you wish to do manual selection for remote devices, turn Automatic Selection off.
16. Select the thin pool name.
17. Click **Add to Storage Group** checkbox and select a storage group.
18. Click **NEXT** to go to the **Summary** page.
19. Review the changes.
20. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## 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.

To delete SRDF pairs from the SRDF List Volumes View, refer to [Deleting SRDF pairs from the SRDF List Volumes View](#) on page 421.

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 5876 or HYPERMAX OS 5977 or higher.


You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

This procedure supports the deletion of SRDF pairs from a SRDF/Metro group.

To delete SRDF pairs:

### Procedure

1. Select the storage system.
2. Select **Data Protection > SRDF**.
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:
    -  Select a group, click , and select **Delete Pair**.
    - Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - Select the **Half Delete** option if deleting one side of the volume pair.
    - 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.
    - Only one side of the RDF device pairs that are removed from the SRDF/Metro session will remain host-accessible when the operation completes. The **Keep R1** or **Keep R2** option is used to specify the side that should remain host-accessible.
    - Click **Advanced Options**. Select the advanced [SRDF session options](#) and click **OK**.



- 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
- Pair level:
  - Select a group and click .
  - Select one or more pairs and click **Delete Pair**.
  - Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
  - Select the **Half Delete** option if deleting one side of the volume pair.
  - 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.
  - Click **Advanced Options**. Select the advanced [SRDF session options](#) and click **OK**.
  - 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## 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. This procedure supports moving SRDF pairs to a SRDF/Metro group.

To move SRDF pairs:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
3.
  - 
  - Select a group, click , and select **Move**.
4. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
5. Select **New SRDF Group**.



6. Select **Full Move** or **Half Move**.
7. Optional: Select **Use Consistency Exempt**.  
This allows volumes to be added, removed, or suspended without affecting the state of the SRDF/A session.
8. Only one side of the RDF device pairs that are moved from the SRDF/Metro session will remain host-accessible when the operation completes. The **Keep R1** or **Keep R2** option is used to specify the side that should remain host-accessible.
9. Click **Advanced Options**. Select the advanced [SRDF session options](#) and click **OK**.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## 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 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.

It is not allowed to set SRDF devices in the non-Metro SRDF mirror to operate in Synchronous mode.

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.



To set SRDF mode:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 , and select **Set Mode**.

- Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
- Select **SRDF Mode**, **Adaptive Copy Mode** and **AC Skew** to set the type of [SRDF session modes](#).
- Select **Use Consistent** to set consistent transition from asynchronous to synchronous mode.
- Pair level:
  - Select a group, click  and click the number next to SRDF Pairs.
  - Select one or more pairs, click , and select **Set Mode**.
  - Select **SRDF Mode**, **Adaptive Copy Mode** and **AC Skew** to set the type of [SRDF session modes](#).
  - Select **Use Consistent** to set consistent transition from asynchronous to synchronous mode.
- 4. Click **Advanced Options**. Select the advanced [SRDF session options](#) and click **OK**.
- 5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing SRDF volume pairs

This procedure explains how to view and manage the volume pairs in a SRDF group.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF**.
3. Select a device group from the list and click  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:

 —[Viewing SRDF volume pair details](#) on page 412

**Establish**—[Establishing SRDF pairs](#) on page 421

**Split**—[Splitting SRDF pairs](#) on page 436

**Suspend**—[Suspending SRDF pairs](#) on page 436

**Restore**—[Restoring SRDF pairs](#) on page 433

**Resume**—[Resuming SRDF links](#) on page 429

**Failover**—[Failing over](#) on page 422

**Failback**—[Failing back](#) on page 423

**Set SRDF/A**—[Setting SRDF/A controls to prevent cache overflow](#) on page 431



**Invalidate**—[Invalidating R1/R2 volumes](#) on page 424

**Ready**—[Making R1/R2 volumes ready](#) on page 425

- Not Ready**—[Making R1/R2 volumes not ready](#) on page 426
- R1 Update**—[Updating R1 volumes](#) on page 438
- RW Enable**—[Read/write enabling R1/R2 volumes](#) on page 428
- Write Disable**—[Read/write disabling R1/R2 volumes](#) on page 429
- RW Disable R2**—[Read/write disabling R2 volumes](#) on page 427
- Refresh**—[Refreshing R1 or R2 volumes](#) on page 430
- Set Mode**—[Setting SRDF mode](#) on page 409

## Viewing SRDF volume pair details

### Procedure

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  to open the SRDF Pair List view.
4. On the Local tab, select the pair and click  to open its details view.

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.

There are 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 the number next to **SRDF Group** will open a view listing the volumes contained in the SRDF group.

## Viewing SRDF volume pair details

This procedure explains how to view an SRDF pair's SRDF group.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF Groups**.
3. Select a device group from the list and click **SRDF Groups** to open the **SRDF Pair List** view.
4. On the **Local** tab, select the pair and click **SRDF Groups** to open its details view.
5. Click the number next to **SRDF Group** to open the Pair's **SRDF Group** view.

The following properties display:

**Group**—RDF group number.

**SRDF Group Label**—RDF group label.

**Remote SRDF Group**—Remote SRDF Group ID.

**Remote Symmetrix**—Remote Symmetrix ID.

**SRDF Group Flags**—SRDF group flags.

**Volume Count**—Number of volumes in the group.

**Copy Jobs**—Maximum number of RDF copy jobs per RDF group.

**Link Limbo (sec)**—Number of seconds (0-10) for the Symmetrix system to continue checking the local RDF link status.

**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

**Minimum Cycle Time**—Minimum cycle time (seconds) configured for this session.

**Session Priority**—Priority used to determine which SRDF/A sessions to drop if cache becomes full. Values range from 1 to 64, with 1 being the highest priority (last to be dropped).


**Transmit Idle Time**—Whether SRDF/A Transmit Idle state is active for the RDF group.

## Viewing SRDF protected storage group pairs

The SRDF SG pair list displays a notification if a capacity mismatch exists between R1 and R2 devices. Mismatch can be  $R1 > R2$  or  $R1 < R2$

To view SRDF group volumes, refer to [Viewing SRDF group volumes](#) on page 454.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups**.
3. Click **SRDF**.
4. Select a storage group instance and click  to open the Storage Group pair list view.
5. Click the number next to **SRDF pairs** to open the SRDF pair list view.

Two tabs are displayed : **Local** and **Hop2**. The non-metro leg of a concurrent RDF pair is viewable in the SRDF/Metro view and the SRDF/Metro leg of the concurrent RDF pair is viewable in the standard RDF view.

The following properties display in the **Local** tab:

**Source Volume**—The name of the source volume.

**Source Type**—The source type of the source volume.

**SRDF Group**—RDF group number.

**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.

The following properties display in the **Hop2** tab:

**Concurrent Volume**—The name of the concurrent volume.

**Symmetrix ID**—Storage system ID.

**SRDF Group**—RDF group number.

**Remote Symmetrix**—Remote Symmetrix ID.

**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.

**SRDF Mode**—The SRDF copy mode.

The following controls are available, depending on the operating environment:



---

**Note**

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

---

**Note**

The dialogs associated with controls listed below do not display the **Use 2nd Hop** option if the hop2 is SRDF/Metro.

---

**Note**

In the event of a concurrent SRDF SG where one leg is SRDF/Metro and one is not SRDF/Metro, the action launching the dialog (Metro or non-Metro) preselects the correct RDFG in the combination box and disables edits on it. The selected RDFG is the one for the SRDF mode of the launching SG.

---

 — [Viewing SRDF protected storage group pair properties](#) on page 417

**Establish**—[Establishing SRDF pairs](#) on page 421

**Split**—[Splitting SRDF pairs](#) on page 436

**Suspend**—[Suspending SRDF pairs](#) on page 436

**Restore**—[Restoring SRDF pairs](#) on page 433

**Resume**—[Resuming SRDF links](#) on page 429

**Delete Pair**—[Deleting SRDF pairs](#) on page 407

**Move**—[Moving SRDF pairs](#) on page 408

**Set Mode**—[Setting SRDF mode](#) on page 409

**Set Volume Attributes > Invalidate**—[Invalidating R1/R2 volumes](#) on page 424

**Set Volume Attributes > Ready**—[Making R1/R2 volumes ready](#) on page 425

**Set Volume Attributes > R1 Update**—[Updating R1 volumes](#) on page 438

**Set Volume Attributes > RW Enable**—[Read/write enabling R1/R2 volumes](#) on page 428

**Set Volume Attributes > Write Disable**—[Read/write disabling R1/R2 volumes](#) on page 429

**Set Volume Attributes > RW Disable R2**—[Read/write disabling R2 volumes](#) on page 427



**Set Volume Attributes > Refresh**—[Refreshing R1 or R2 volumes](#) on page 430

**Set SRDF/A**—[Setting SRDF/A controls to prevent cache overflow](#) on page 431

## Viewing SRDF protected storage group pair properties

### Procedure

1. Select the storage system.

2. Select **DATA PROTECTION > Storage Groups**.
3. Click **SRDF**.
4. Select a storage group and click  to open the storage group list view.
5. Click the number next to **SRDF pairs** to open the SRDF pair list view.
6. Select a pair and click  to open the SRDF pair list properties panel.

The following properties display, depending on the operating environment:

**Storage Group**—The storage group ID.

**Local Volume**—The local volume ID.

**SRDF Group Number**—SRDF group number.

**Remote SRDF Group Number**—Remote SRDF group number.

**Remote Volume**—The remote volume ID.

**Pair State**—The state of the SRDF pair.

**SRDF Mode**—The SRDF mode.

**Adaptive Copy Mode**—The adaptive copy mode.

**Adaptive Copy Skew**—The adaptive copy skew.

**Consistency State**—The consistency state.

**Consistency Exempt**—Indicates consistency exempt status.

**Link Status**—Indicates link state.

**Link Domino** —Indicates link Domino state.

**Local Volume Invalid R1 Track Count** —Indicates Local Volume Invalid R1 Track Count.

**Local Volume Invalid R2 Track Count** —Indicates Local Volume Invalid R2 Track Count.

**Local Volume SRDF State** —Indicates SRDF state of the local volume.

**Local Volume SRDF Type**—Indicates SRDF type of the local volume.

**Local Volume Remote Write Pacing Track Count**—Indicates Local Volume Remote Write Pacing Track Count.

**Local Volume Track Size**—Indicates track size of the local volume.

**Remote Local Volume Invalid R1 Track Count** —Indicates Remote Volume Invalid R1 Track Count.

**Remote Volume Invalid R2 Track Count** —Indicates Remote Volume Invalid R2 Track Count.

**Remote Volume SRDF State** —Indicates SRDF state of the remote volume.

**Remote Volume Remote Write Pacing Track Count**—Indicates Remote Volume Remote Write Pacing Track Count.

**Remote Volume Track Size**—Indicates track size of the remote volume.

**SRDF/A Pacing capable**—Indicates SRDF/A pacing capability.

**Configured Group Level Exempt State**—Configured Group Level Exempt state indication.

**Effective Group Level Exempt State**—Effective Group Level Exempt state indication.

**Volume Level Pacing State**—Volume Level Pacing state indication.

**SRDF/A Consistency Protection**—SRDF/A Consistency Protection indication.

**SRDF/A Average Cycle Time**—SRDF/A Average Cycle Time.

**SRDF/A Minimum Cycle Time**—SRDF/A Minimum Cycle Time.

**SRDF/A Cycle Number**—SRDF/A Cycle Number.

**SRDF/A Session Number**—SRDF/A Session Number.

**Transmit Queue Depth of R1 side**—Transmit Queue Depth of R1 side.

**SRDF/A Uncommitted Tracks Count**—SRDF/A Uncommitted Tracks count.

**SRDF/A Number of Volumes in Session**—SRDF/A Number of Volumes in session.

**SRDF/A Session Uncommitted Tracks Count**—SRDF/A Session Uncommitted Tracks count.

**SRDF/A R1 DSE Used Track Count**—SRDF/A R1 DSE Used Track c.

**SRDF/A R1 Cache In Use Percent**—SRDF/A R1 Cache In Use Percent.

**SRDF/A R1 Shared Track Count**—SRDF/A R1 Shared Track count.

**SRDF/A R1 to R2 Lag Time**—SRDF/A R1 to R2 Lag Time.

**SRDF/A R2 DSE Used Track Count**—SRDF/A R2 DSE Used Track count.

**SRDF/A R2 Cache In Use Percent**—SRDF/A R2 Cache In Use Percent.

**SRDF/A Session Minimum Cycle Time**—SRDF/A Session Minimum Cycle time.

**SRDF/A Transmit Idle State**—SRDF/A Transmit Idle state.

**SRDF/A Transmit Idle Time**—SRDF/A Transmit Idle time.

**Suspended State**—Suspended state.

**SQAR Mode**—SQAR Mode status (enabled or disabled).

There are also links to views displaying objects contained in and associated with the SRDF pair. Each link is followed by a number, indicating the number of objects in the corresponding view. For example, clicking the number next to **SRDF Group Number** opens a view listing the related SRDF groups.

## 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.

To delete SRDF pairs from the SRDF List Volumes View, refer to [Deleting SRDF pairs from the SRDF List Volumes View](#) on page 421.

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 5876 or HYPERMAX OS 5977 or higher.

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

This procedure supports the deletion of SRDF pairs from a SRDF/Metro group.

To delete SRDF pairs:

### Procedure


1. Select the storage system.
2. Select **Data Protection > SRDF**.
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:
    -  Select a group, click , and select **Delete Pair**.
    - Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - Select the **Half Delete** option if deleting one side of the volume pair.
    - 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.
    - Only one side of the RDF device pairs that are removed from the SRDF/Metro session will remain host-accessible when the operation completes. The **Keep R1** or **Keep R2** option is used to specify the side that should remain host-accessible.
    - Click **Advanced Options**. Select the advanced [SRDF session options](#) and click **OK**.
    - 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 920 and [Previewing jobs](#) on page 920.
      - Expand **Add to Job List**, and click **Run Now** to perform the operation now.
  - Pair level:
    -  Select a group and click .
    - Select one or more pairs and click **Delete Pair**.
    - Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
    - Select the **Half Delete** option if deleting one side of the volume pair.

- 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.
- Click **Advanced Options** . Select the advanced [SRDF session options](#) and click **OK**.
- 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Deleting SRDF pairs from the SRDF List Volumes View

To delete SRDF pairs from the SRDF List Volumes View:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF Groups**
3. Select a group and click .
4. Click the number next to **Volumes**.
5. Select a volume and click and select **Delete Pairs** to open the **Delete Pair** dialog box.
6. Select the **Half Delete** option if deleting one side of the volume pair.
7. 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.
8. Optional: Select **Use Force**.
9. Click **OK**.

## Establishing SRDF pairs

### Before you begin


SRDF requires Engenuity version 5876 or HYPERMAX OS 5977 or higher.

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

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:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 **Establish**.
    - b. Select **Full** or **Incremental** session type.
    - c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - 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.
  - Pair level:
    - a. Select a group, click , and click the number next to SRDF Pairs.
    - b. Select one or more pairs and click **Establish**.
    - c. Select **Full** or **Incremental** establish type.
    - d. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - e. 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.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Failing over

### 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.







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.

The Failing Over operation is not allowed on the non-Metro SRDF mirror if it results in a Metro device becoming write-disabled (WD).

To initiate a failover:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF or DATA PROTECTION > Device Groups > SRDF**.
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 , and select **Failover** .
    - 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.
  - Pair level:
    - a.
 
  
Select a group, click , and click the number next to SRDF Pairs.
    - b.
 
  
Select one or more pairs, click , and select **Failover** .
    - c. Select the fail over.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Failing back

### Before you begin






SRDF requires Enginuity version 5876 or HYPERMAX OS 5977 or higher.

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.

To initiate a failback:

### Procedure

1. Select the storage system.

2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 , and select **Failback**.
    - 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.
  - Pair level:
    - a. Select a group, click , and click the number next to SRDF Pairs.
    - b.  Select one or more pairs, click , and select **Failback**.
    - c. Select the fail over.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Invalidating R1/R2 volumes

### Before you begin

SRDF requires Enginuity version 5876 or HYPERMAX OS 5977 or higher.

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.


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.



To invalidate R1/R2 volumes:



### Procedure


1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 , and select **Set Volume Attributes > Invalidate**.
  - b. Select R1 or R2 volume type.
  - c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
- Pair level:
  - a.
 

Select a group, click , and click the number next to SRDF Pairs.
  - b.
 

Select one or more pairs, click , and select **Set Volume Attributes > Invalidate**.
  - c. Select side **R1** or **R2**.
- 4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
- 5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Making R1/R2 volumes ready


### Before you begin


SRDF requires Enginuity version 5876 or HYPERMAX OS 5977 or higher.



You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

To make R1 or R2 volumes ready to their local hosts:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 , and select **Set Volume Attributes > Ready**.
    - b. Select side **R1** or **R2**.
    - c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).

- Pair level:
  - a. Select a group, click , and click the number next to SRDF Pairs.
  - b.
    - Select one or more pairs, click , and select **Set Volume Attributes > Ready**.
  - c. Select **R1** or **R2** volume type.
- 4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
- 5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Making R1/R2 volumes not ready

### Before you begin



SRDF requires Enginuity version 5876 or HYPERMAX OS 5977 or higher.


You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

This procedure explains how to set the source (R1) or the target (R2) volumes not ready to the local host.

To make R1/R2 volumes not ready:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 , and select **Set Volume Attributes > Not Ready**.
    - b. Select side **R1** or **R2**.
    - c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
  - Pair level:
    - a. Select a group, click , and click the number next to SRDF Pairs.

- b.
  - Select one or more pairs, click , and select **Set Volume Attributes > Not Ready**.
- c. Select **R1** or **R2** volume type.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Read/write disabling R2 volumes




### Before you begin

SRDF requires Enginuity version 5876 or HYPERMAX OS 5977 or higher.

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

To read/write disable R2 volumes:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 , and select **Set Volume Attributes > RW Disable R2**.
    - b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
  - Pair level:
    - a.
      - Select a group, click , and click the number next to SRDF Pairs.
    - b.
      - Select one or more pairs, click , and select **Set Volume Attributes > RW Disable R2**.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:

- Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Read/write enabling R1/R2 volumes




### Before you begin

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

This procedure explains how to write enable the R1 (source) or R2 (target) volumes ready to their local hosts.

To read/write enable R1/R2 volumes:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF or DATA PROTECTION > Device Groups > SRDF**.
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 , and select **RW Enable**.
    - b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - c. Select **RW Enable R1s** or **RW Enable R2s** volume type.
  - Pair level:
    - a. Select a group and click , and click the number next to SRDF Pairs.
    - b. Select one or more pairs, click , and select **RW Enable**.
    - c. Select **R1** or **R2** volume type.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Resuming SRDF links



### Before you begin

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

This procedure explains how to resume I/O traffic on the SRDF links for all remotely mirrored SRDF pairs in the group.

To resume SRDF links:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
3.  Select a group, click , and select **Resume**.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Read/write disabling R1/R2 volumes

### Before you begin

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.




This procedure explains how to write disable source (R1) volumes/target (R2) volumes to their local hosts.

The Write Disable R1 operation is not allowed on the non-Metro RDF mirror if it results in a Metro device becoming write-disabled (WD).

To write disable R1/R2 volumes:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups** and click the **SRDF** tab.
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 , and click **Write Disable**.
  - Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
  - Select Write Disable R1s or Write Disable R2s volume type.
- Pair level:
  - Select a group and click 
  - - Select one or more pairs, click , and select **Write Disable**.
    - Select R1 or R2 volume type.
- 4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
- 5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Refreshing R1 or R2 volumes

### Before you begin

To invoke this operation, the SRDF pair(s) must be in one of the following states:

- Suspended and Write Disabled at the source
- Suspended and Not Ready at the source
- Failed Over with the -force option specified
- This operation is rejected if the target has invalid local (R2) tracks.


You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.


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.



To refresh volumes:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF or DATA PROTECTION > Device Groups > SRDF**.
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 , and select **Set Volume Attributes > Refresh**.
- b. Select R1 or R2 volume type.
- c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
- Pair level:
 

This action can also be run from pair level details view. Select a pair and click .

  - a. Select a group and click , and click the number next to SRDF Pairs.
  - b.
    - Select one or more pairs, click , and select **Set Volume Attributes > Refresh**.
  - c. Select R1 or R2 volume type.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting SRDF/A controls to prevent cache overflow

### Before you begin


SRDF requires Enginuity version 5876 or HYPERMAX OS 5977 or higher.

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

This procedure explains how to activate or deactivate SRDF/A control actions 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.

To activate or deactivate SRDF/A controls:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
3.
  - Select a group, click more , and select **Asynchronous > Set SRDF/A**.
4. Select **Activate SRDF/A** or **Deactivate SRDF/A**.
5. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).

6. Select **Activate Type** or **Deactivate Type**.
7. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
8. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.


## Setting consistency protection


### Before you begin

To set consistency protection:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF or DATA PROTECTION > Device Groups > SRDF**.
3.
 



 Select a group, click more , and select **Asynchronous > Set Consistency**.
4. select **Enable** or **Disable**.
5. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
6. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
7. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.


## 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.

To reset the original device identity:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF groups**.
3. Do the following, depending on whether you want to perform the operation at the group level or pair level:
  - Group level
 



 Select a former unbiased SRDF/Metro storage group and click **Reset SRDF/Metro Identity** to open the **Reset Original Identity** dialog box.



- Pair level:
  - Select **STORAGE > Volumes**.
  - Filter the view to display volume(s) that were formally part of a SRDF/Metro pair.
  - Do one of the following:
    - Select a volume and click **Reset SRDF/Metro Identity**.
    - Select a volume, click , and then click **Reset SRDF/Metro Identity**.
- 4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
- 5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## 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 5876 or HYPERMAX OS 5977 or higher.

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

To restore SRDF pairs:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 **Restore**.
    - b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - 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.

- Pair level:
  - a. Select a group, click  and click the number next to SRDF Pairs.
  - b. Select one or more pairs and click **Restore**.
  - c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
  - d. Select **Full** or **Incremental** restore type.
  - e. 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.
- 4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
- 5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## 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.

---

### Note


**Set Bias** cannot be invoked for a witness protected SRDF/Metro group.


---

To set bias:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
3.
 



 Select a group, click , and select **Set Bias**.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting the SRDF GCM flag

This procedure supports the setting of the SRDF GCM flag at the Storage Group level and at the individual volume level.

The Geometry Compatible Mode (GCM) parameter modifies how a storage system running HYPERMAX OS 5997 or later manages the size of a volume. When the GCM attribute is set, the volume is treated as  $\frac{1}{2}$  a cylinder smaller than its true configured size. This enables a volume on a storage system running HYPERMAX OS 5977 to be paired with a volume on an storage system running Enginuity 5876, when the 5876 volume has an odd number of cylinders.

Before you begin:

SRDF requires HYPERMAX OS 5977 or later.

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF groups**.
3. Do the following, depending on whether you want to perform the operation at the group level or volume level:

Group level:

- a.
  - Select a storage group and click **Set SRDF GCM** to open the **Set GCM** dialog box.
  - Click **On** to set the GCM flag or **Off** to unset the flag.

---

#### Note

The only way to unset this flag is to unmap the device which requires an outage at the host which would mean losing access to volumes.

---

- 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 **On** to set the GCM flag or **Off** to unset the flag.

---

#### Note

The only way to unset this flag is to unmap the device which requires an outage at the host which would mean losing access to volumes.

---

- 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 later.

To set the volume state:

#### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**.

3. Select a former unbiased SRDF/Metro storage group and click **Set Volume Status**.
4. Click **OK**.

## Splitting SRDF pairs

This procedure explains how to stop SRDF pair mirroring.


### Before you begin

SRDF requires Enginuity version 5876 or HYPERMAX OS 5977 or higher.

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

To split SRDF pairs:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 **Split**.
    - b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - c. Select **Use Immediate** for immediate split on asynchronous devices.
  - Pair level:
    - a. Select a group, click  and click the number next to SRDF Pairs.
    - b. Select one or more pairs and click **Split**.
    - c. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - d. Select **Use Immediate** for immediate split on asynchronous devices.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Suspending SRDF pairs


This procedure explains how to stop data transfer between SRDF pairs.

Before you begin:

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard RDF view. If you are viewing a storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

To suspend SRDF pairs:

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 **Suspend**.
    - Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - Select **Use Immediate** or **Use Consistency Exempt**
    - 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. This option is not allowed until all the devices in the SRDF/Metro config, both new and existing, are in the **ActiveActive** or **ActiveBias** SRDF pair state.
  - Pair level:
    - Select a group, click  and click the number next to SRDF Pairs.
    - Select one or more pairs and click **Suspend**.
    - Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - Select **Use Immediate** or **Use Consistency Exempt**
    - 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. This option is not allowed until all the devices in the SRDF/Metro config, both new and existing, are in the **ActiveActive** or **ActiveBias** SRDF pair state.
    - Only one side of the RDF device pairs that are suspended from the SRDF/Metro session will remain host-accessible when the operation completes. The **Keep R1** or **Keep R2** option is used to specify the side that should remain host-accessible. This applies to storage systems running PowerMaxOS 5978 only.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## 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 5876 or HYPERMAX OS 5977 or higher.

To swap SRDF personalities:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF or DATA PROTECTION > Device Groups > SRDF**.
3.
 



 Select a group, click , and select **Swap**.
4. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
5. For optional [Refreshing R1 or R2 volumes](#) on page 430, select **R1, R2 or None**.
6. For optional half swapping, select **Half Swap**.

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).

7. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
8. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## 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 5876 or HYPERMAX OS 5977 or higher.

You are not able to perform SRDF/Metro control actions at the SG level on the SRDF/Metro pairs in a standard SRDF view and you are not allowed to perform standard SRDF actions on the SRDF/Metro leg in a standard SRDF view. If you are viewing a

storage system not associated with either side of the pair of interest then you need to go to the view of the relevant storage system.

To update R1 volumes:

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF** or **DATA PROTECTION > Device Groups > SRDF**.
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 , and click **Set Volume Attributes > R1 Update**.
    - b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (not applicable if the hop2 is SRDF/Metro).
    - c. Select **Remote** if R1 volumes are a remote.
  - Pair level:
    - a. Select a group, click  and click the number next to SRDF Pairs.
    - b.  Select one or more pairs, click , and select **Set Volume Attributes > R1 Update**.
    - c. Select **Remote** if R1 volumes are a remote.
4. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
5. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## 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

Session option	Description	Available with action
		Write Disable R1 Ready R1 Ready R2 RWDisableR2 Enable Disable
Consistent	Allows only consistent transition from async to sync mode.	Activate
Consistency Exempt	Allows you to add or remove volumes from an RDF group that is in Async mode without requiring other volumes in the group to be suspended.	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	



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

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

## 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

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

## RBAC roles for performing local and remote replication actions

The table below details the roles needed to perform SRDF local and remote replication actions.

### Note

Unisphere for PowerMax does not support RBAC Device Group management.

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

## Understanding Virtual Witness

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

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

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

The following Virtual Witness tasks can be performed from Unisphere.

[Viewing Virtual Witness instances](#)

[Adding a Virtual Witness](#)

[Viewing Virtual Witness instance details](#)

[Enabling a Virtual Witness](#)

[Disabling a Virtual Witness](#)

[Removing Virtual Witness](#)

## Adding SRDF Virtual Witness instances

### Before you begin

Unisphere provides monitoring and management for SRDF/Metro Virtual Witness instances on Virtual Witness capable storage systems running HYPERMAX OS 5977 Q3 2016 or higher.

A Virtual Witness instance needs to be created for both participating arrays.

See [Understanding Virtual Witness](#) on page 444 for additional information.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Virtual Witness** to open the **Virtual Witness** list view.
3. Click **Create**.
4. Type values for the following:
  - **Virtual Witness Name**—User-defined Virtual Witness instance name.
  - **IP/DNS**—IPv4 or IPv6 address, or DNS name from embedded Guest that is associated with Virtual Witness instance.
5. Optional: Select the **Add Virtual Witness to remote arrays** checkbox and select the arrays (these arrays support the Virtual Witness functionality) that are to have the same Virtual Witness added.
- 6.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Removing SRDF Virtual Witness instances

### Before you begin

Unisphere provides monitoring and management for SRDF/Metro Virtual Witness instances on Virtual Witness capable storage systems running HYPERMAX OS 5977 Q3 2016 or higher.

You cannot remove a Virtual Witness instance that is in use (protecting one or more SRDF/Metro sessions).

See [Understanding Virtual Witness](#) on page 444 for additional information.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Virtual Witness** to open the **Virtual Witness** list view.
3. Select a virtual witness instance and click **DELETE**.
4. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Set state for SRDF Virtual Witness instances

### Before you begin

Unisphere provides monitoring and management for SRDF/Metro Virtual Witness instances on Virtual Witness capable storage systems running HYPERMAX OS 5977 Q3 2016 or higher.

The Virtual Witness disable operation may or may not require additional force flags based on if it is currently protecting SRDF/Metro sessions and if an alternate witness is available. If the vWitness is currently protecting Metro Sessions, the storage system performs a search for replacement Witnesses (virtual or physical) to use.

You cannot disable a Virtual Witness instance that is in use (protecting one or more SRDF/Metro sessions).

The **Set State** operation changes the state of the Virtual Witness instance from enabled to disabled or from disabled to enabled.

See [Understanding Virtual Witness](#) on page 444 for additional information.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Virtual Witness** to open the **Virtual Witness** list view.
3. Select a Virtual Witness instance and click **Set State**.

Note: When disabling an enabled Virtual Witness instance:

- Click **Advanced Options** and select the **Use Force** check box.  
The command fails if the virtual Witness is currently in use (protecting a SRDF/Metro Session) and there is another witness (either virtual or physical) that is available to take over for it. The `force` flag is needed in order to continue.
  - Click **Advanced Options** and select the **Use SymForce** check box.  
The command fails if the virtual Witness is currently in use (protecting a SRDF/Metro Session) and there is no other witness (either virtual or physical) that is available to take over for it. The `symforce` flag is needed in order to continue
4. Do one of the following:
    - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.

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

## Viewing SRDF Virtual Witness instances

### Before you begin

Unisphere provides monitoring and management for SRDF/Metro Virtual Witness instances on Virtual Witness capable storage systems running HYPERMAX OS 5977 Q3 2016 or higher.

A Virtual Witness needs to be created for both participating arrays.

See [Understanding Virtual Witness](#) on page 444 for additional information.


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Virtual Witness** to open the **Virtual Witness** list view.

The following properties display, depending on the operating environment:

- **Witness name**—User-defined Virtual Witness instance name.
- **State**—State of Virtual Witness instance .
- **Alive**—Flag to indicate if the Virtual Witness instance is alive.
- **In Use**—Flag to indicate if the Virtual Witness instance is in use.

The following controls are available, depending on the operating environment:

-  —[Viewing SRDF Virtual Witnesses details](#) on page 447
- **Create**—[Adding SRDF Virtual Witness instances](#) on page 445
- **Set State**—[Set state for SRDF Virtual Witness instances](#) on page 446
- **Delete**—[Removing SRDF Virtual Witness instances](#) on page 445


## Viewing SRDF Virtual Witnesses details

### Before you begin

Unisphere provides monitoring and management for SRDF/Metro Virtual Witness on Virtual Witness capable storage systems running HYPERMAX OS 5977 Q3 2016 or higher.

See [Understanding Virtual Witness](#) on page 444 for additional information.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Virtual Witness** to open the **Virtual Witness** list view.
3. Select a Virtual Witness instance and click  to open the **Details** view.

The following properties display:

- **Witness name**—User-defined witness name.

- **IP/DNS**—IPv4 or IPv6 address, or DNS name from embedded Guest that is associated with Virtual Witness instance. .
- **Port**—Port associated with Virtual Witness instance.
- **Alive**—Flag to indicate if the Virtual Witness instance is alive.
- **State**—State of Virtual Witness instance.
- **InUse**—Flag to indicate if the Virtual Witness instance is in use.
- **Duplicate**—Flag to indicate if the Virtual Witness instance is a duplicate. A duplicate witness is a witness which has the same unique ID as another witness on the storage system, for example, in the case where it was added twice.
- **SRDF Groups**—Number of SRDF groups.

There are links to views for objects associated with the Virtual Witness instance. 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 **SRDF Groups** opens the view listing the SRDF Groups associated with the Virtual Witness instance.

## Creating SRDF/A DSE pools

### Before you begin

SRDF/A DSE pools are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3. Click **Create**.

You can also create DSE pools from the DSE pools 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.
6. Optional: Click on slider bar to enable the new pool member(s).
7. Click **OK**.

## Deleting SRDF/A DSE pools

### Before you begin

SRDF/A DSE pools are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3. Select a pool and click  .



4. Click **OK**.

## Adding volumes to SRDF/A DSE pools

### Before you begin

SRDF/A DSE pools are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3. Select a pool and click **Add**.
4. Select the volumes to add.
5. Optional: Click on slider bar to enable the new pool member(s).
6. Click **OK**.

## Removing volumes from SRDF/A DSE pools

### Before you begin

SRDF/A DSE pools are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3. Select a pool and click **Remove**.
4. Click **OK**.



## Enabling all volumes in SRDF/A DSE pools

### Before you begin

SRDF/A DSE pools are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3.
 


 Select a pool, click  , and select **Enable All**.
4. Click **OK**.

## Disabling all volumes in SRDF/A DSE pools


### Before you begin


SRDF/A DSE pools are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.

2. Select **DATA PROTECTION > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.
3.
 



 Select a pool, click , and select **Disable All**.
4. Click **OK**.

## Viewing SRDF/A DSE pools

### Before you begin

SRDF/A DSE pools are supported on storage systems running Enginuity 5876.

### Procedure



1. Select the storage system.
2. Select **DATA PROTECTION > 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 storage system.

The following properties display:

- **Name**—Name of the pool.
- **DSE 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.

The following controls are available:

-  — [Viewing SRDF DSE pool details](#) on page 450
- **Create**—[Creating SRDF/A DSE pools](#) on page 448
- **Add**—[Adding volumes to SRDF/A DSE pools](#) on page 449
-  — [Deleting SRDF/A DSE pools](#) on page 448
- **Enable All**—[Enabling all volumes in SRDF/A DSE pools](#) on page 449
- **Disable All**—[Disabling all volumes in SRDF/A DSE pools](#) on page 449
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945

## Viewing SRDF DSE pool details

### Before you begin

SRDF/A DSE pools are supported on storage systems running Enginuity 5876.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF/A DSE Pools** to open the **SRDF/A DSE Pools** list view.

3. Select the pool and click  to open its **Details** view.

Use the SRDF/A DSE Pool **Details** view to display and manage a TimeFinder/Snap pool.

The following properties display:

- **Array ID**—Storage system on which the pool resides.
- **DSE Pool Name**—Name of the pool.
- **Pool Type**—Pool type.
- **Emulation**—Emulation type.
- **RAID Protection**—Protection level of the volumes in the pool.
- **Technology**—Technology on which the volumes in the pool reside.
- **Pool State**—Whether the pool is Enabled or Disabled.
- **Num 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.
- **% Used**—Percentage used in GB.
- **Used (GB)**—Total used in GB.
- **Free (GB)**—Total free space in GB.

The properties 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 the number next to **NumVolumes** opens a view listing the SAVE volumes contained in the pool.

## Creating TimeFinder/Snap pools

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Click **Create**.
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.
6. Optional: To enable new volumes in the pool, select **Enable new pool Member**.  
The total enabled pool capacity in GB is displayed.
7. Click **OK**.

## Adding volumes to TimeFinder/Snap pools

TimeFinder/Snap pools are supported on storage systems running Enginuity OS 5876.


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Select a pool and click  to open its **Details** view.
4. Click the number next to **Num Volumes** to open the **SAVE Volumes** view.
5. Select one or more volumes and click **Add**.
6. Select one or more volumes.
7. Optional: To enable new volumes in the pool, select **Enable new pool Member**.  
The total enabled pool capacity in GB is displayed.
8. Click **OK**.

## Enabling all volumes in TimeFinder/Snap pools

TimeFinder/Snap pools are supported on storage systems running Enginuity OS 5876.


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > TimeFinder Snap Pools**.
3. Select a snap pool, click , and select **Enable All**.
4. Click **OK**.

## Disabling all volumes in TimeFinder/Snap pools

TimeFinder/Snap pools are supported on storage systems running Enginuity OS 5876.

### Procedure


1. Select the storage system.
2. Select **DATA PROTECTION > TimeFinder Snap Pools**.
3. Select a snap pool, click , and select **Disable All**.
4. Click **OK**.

## Deleting TimeFinder/Snap Pools

TimeFinder/Snap pools are supported on storage systems running Enginuity OS 5876.

### Procedure


1. Select the storage system.
2. Select **DATA PROTECTION > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.

3. Select a pool and click .
4. Click **OK**.

## Removing volumes from TimeFinder/Snap pools

TimeFinder/Snap pools are supported on storage systems running Enginuity OS 5876.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Select a pool and click  to open its **Details** view.
4. Click the number next to **Num Volumes** to open the **SAVE Volumes** view.
5. Select one or more volumes and click **Remove**.
6. Click **OK**.

## Viewing TimeFinder/Snap pools

TimeFinder/Snap pools are supported on storage systems running Enginuity OS 5876.

### Procedure



1. Select the storage system.
2. Select **DATA PROTECTION > 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 storage 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.
- **Capacity (GB)**—Capacity in GB.
- **Pool State**—Whether the pool is Enabled or Disabled.
- **% Used**—Percentage of pool used.
- **Used (GB)**—Total used space in GB.
- **Free (GB)**—Total free space in GB.

The following controls are available:


-  —[Viewing TimeFinder/Snap pool details](#) on page 454
- **Create**—[Creating TimeFinder/Snap pools](#) on page 451
- **Add**—[Adding volumes to TimeFinder/Snap pools](#) on page 452
-  —[Deleting TimeFinder/Snap Pools](#) on page 452

- **Enable All**—[Enabling all volumes in TimeFinder/Snap pools](#) on page 452
- **Disable All**—[Disabling all volumes in TimeFinder/Snap pools](#) on page 452
- **Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945

## Viewing TimeFinder/Snap pool details

TimeFinder/Snap pools are supported on storage systems running Enginuity OS 5876.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > TimeFinder Snap Pools** to open the **TimeFinder Snap Pools** list view.
3. Select a pool and click  to open its **Details** view.

The following properties display:

- **Array ID** —Storage system on which the pool resides.
- **Name** —Name of the pool.
- **Pool Type**—Pool type.
- **RAID Protection**—Protection level of the volumes in the pool.
- **Technology**— Technology on which the volumes in the pool reside.
- **Pool State**— State of the pool (Enabled or Disabled).
- **Num 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 (GB)**— Total free space in GB.
- **Used (GB)**—Total used space in GB.

There are 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 **NumVolumes** opens a view listing the SAVE volumes contained in the TimeFinder snap pool.

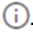
The **Performance Views** panel links you to the performance monitor and analyze views for the snap pool.

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

## Viewing SRDF group volumes

This procedure explains how to view the volumes in an SRDF group:

**Procedure**

1. Select the storage system.
2. Select **Data Protection** > **SRDF groups** to open the **SRDF groups** list view.
3. Select the SRDF group and click  to open its **Details** view.
4. Click the number next to **SRDF Group Volumes** to open the **SRDF List Volumes** view.

The following properties display:

**Volumes**—Local volume ID.

**Configuration**—SRDF configuration.

**Remote Symmetrix**—Remote storage system ID.

**Remote SRDF Group**—Remote SRDF group ID.

**Target Volume**—Target volume ID.

**State**—Session state of the pair

**Pair State**—Volume pair state.

**Remote Volume State**—State of the remote volume.

**SRDF Mode**—SRDF copy type.

## Viewing SRDF protected storage groups

Unisphere 8.1 and higher provides SRDF monitoring and management for storage groups. This includes 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. See [Managing remote replication sessions](#) on page 402 for additional information.

**Procedure**

1. Select the storage system.
2. Select **DATA PROTECTION** > **Storage Groups**.
3. Click the **SRDF** tab.

The following properties display, depending on the operating environment:

- **Storage Group**—User-defined storage group name.
- **States**—The state of the storage group. Possible values are:
  - ActiveActive
  - ActiveBias
  - 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.

- **Modes**—The SRDF modes.
- **SRDF Type**—The SRDF type. SGs with volumes having multiple SRDF types display multiples here, for example, R1 and R2.
- **SRDF Groups**—The SRDF group number. Concurrent SRDF setups list multiple SRDF Group numbers.

Click  to view the following additional properties:

- **Capacity(GB)**—Total capacity of the storage group in GB.
- **SRDF Pairs**—Number of associated SRDF pairs.
- **Masking Views**—The number of associated masking views.
- **Emulation**—The emulation type (ALL, FBA, CKD).
- **Group Type**—The group type.
- **Bias Type**—The bias type.
- **Production Volumes**—The number of production volumes.
- **Last Updated**—The date and time of the last update.

The following controls are available, depending on the operating environment and the mode:



- **Establish**—[Establishing SRDF pairs](#) on page 421
- **Split**—[Splitting SRDF pairs](#) on page 436
- **Suspend**—[Suspending SRDF pairs](#) on page 436
- **Restore**—[Restoring SRDF pairs](#) on page 433
- **Resume**—[Resuming SRDF links](#) on page 429
- **Failover**—[Failing over](#) on page 422
- **Failback**—[Failing back](#) on page 423
- **Swap**—[Swapping SRDF personalities](#) on page 438
- **Move**—[Moving SRDF pairs](#) on page 408
- **Delete Pair**—[Deleting SRDF pairs](#) on page 407
- **Set Mode**—[Setting SRDF mode](#) on page 409
- **Set Bias**—[Setting bias location](#) on page 434
- **Set Volume Attributes > Invalidate**—[Invalidating R1/R2 volumes](#) on page 424
- **Set Volume Attributes > Ready**—[Making R1/R2 volumes ready](#) on page 425
- **Set Volume Attributes > Not Ready**—[Making R1/R2 volumes not ready](#) on page 426



- **Set Volume Attributes > R1 Update**—[Updating R1 volumes](#) on page 438
- **Set Volume Attributes > RW Enable** —[Read/write enabling R1/R2 volumes](#) on page 428
- **Set Volume Attributes > Write Disable**—[Read/write disabling R1/R2 volumes](#) on page 429
- **Set Volume Attributes > RW Disable R2s**—[Read/write disabling R2 volumes](#) on page 427
- **Set Volume Attributes > Refresh**—[Refreshing R1 or R2 volumes](#) on page 430
- **Asynchronous > Set SRDF/A**— [Setting SRDF/A controls to prevent cache overflow](#) on page 431
- **Asynchronous > Set Consistency**—[Setting consistency protection](#) on page 350

## Viewing related SRDF groups

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups**.
3. Click **SRDF**.
4. Select a storage group and click  to open the Storage Group pair list view.
5. Click the number next to **SRDF pairs** to open the SRDF pair list view.
6. Select a pair and click  to open the SRDF pair list properties panel.
7. Click the number next to **SRDF Group Number** to open the related SRDF groups list view.

The following properties display, depending on the operating environment:

**SRDF Group Number**—SRDF group number.

**SRDF Group label**—SRDF group label.

**Remote SRDF Group Number**—Remote SRDF group number.

**Remote Symmetrix ID**—Remote Symmetrix ID.

**Volumes Count** —Indicates Volumes count.

## 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:

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
5876	250	64	64	1 to 250

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

To create an SRDF group:

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF Groups** and click **Create SRDF Group**, or select **CREATE SRDF GROUP** from the **REPLICATION** dashboard.
3. Select a **Communication Protocol** to use when moving data across the SRDF links.

The value you select here will populate the **Director Port** list.

4. Select a **Remote Array ID**.
5. To refresh the remote storage system information, click **Scan**.

The scan operation looks for SRDF capable systems known to Unisphere.

6. Type a **SRDF Group Label** (name).
7. Click **SRDF/Metro Witness Group**.

This checkbox is selectable when the local storage system and the remote selected storage system are both Witness capable.

8. Select a local **SRDF Group Number**.
9. Select the local director ports through which the group will communicate.
10. Click **Advanced Options** to set the advanced options, as described next.

Setting Advanced options:

- a. Select a local **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.
- b. Select (enable) **Local 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.
- c. Select (enable) **Local 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.

- d. Click **OK**.
11. Select a **Remote SRDF Group Number**.
12. Select the remote director ports through which the group will communicate.
13. Click **Advanced Options** to set the advanced options, as described next.

Setting Advanced options:

- a. 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.
- b. 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.
- c. 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.
- d. Click **OK**.
- e. A summary page, displaying all values and options selected, is displayed.
14. Optional: Set one or more of the following:
  - 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.
  - 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 PowerMaxOS 5978 or higher.
15. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Modifying SRDF groups

To modify SRDF groups:

### Procedure


1. Select the storage system.
2. Select **DATA PROTECTION > SRDF Groups** to open the **SRDF groups** list view.
3. Select a group and click **Modify**.

4. Do any number of the following steps:
  - a. Select a new **Local Array** through which the group will communicate. When specifying a local or remote array for a storage system running HYPERMAX OS 5977, you can select one or more SRDF ports.
  - b. Select a new **Remote Array** through which the group will communicate.
5. Select **Advanced Options** and do any number of the following steps:
  - a. Select a local **Link Limbo Period**. The 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.
  - b. 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.
  - c. 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.
  - d. 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 5876 or later.
  - e. 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 5876 or later.
  - f. 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.
  - g. Select (enable) **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.
  - h. Select (enable) **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.
6. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting SRDF/A DSE attributes

### Procedure

1. Select a storage system.
2. Select **DATA PROTECTION > SRDF Groups**.
- 3.

Select a group, click , and select **SRDF/A DSE Setting**.

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.


This feature is supported with thin devices.

8. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting SRDF/A group attributes

### Procedure

1. Select a storage system.
2. Select **DATA PROTECTION > SRDF Groups** to open the **SRDF groups** list view.
- 3.


Select a group, click , and select **SRDF/A Setting**.


4. Type the **Minimum Cycle Time**. This is the minimum amount of time (in seconds) the storage system will wait before attempting to perform an RDF/A cycle switch. Possible values range from 1 to 60 seconds.
5. Type the **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. Select **Transmit Idle Enabled** 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. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Setting SRDF/A pace attributes

### Procedure

1. Select a storage system.
2. Select **DATA PROTECTION > SRDF Groups**.
3.
 



 Select a group, click , and select **SRDF/A Pacing Setting**.
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 Group Pacing** and **Autostart Volume Pacing**). 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. Do one of the following:
  - Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Swapping SRDF groups

### Before you begin

Before you begin:



- If the target (R2) volume is on a storage system running HYPERMAX OS 5977 or later, 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, 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.

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.

To swap SRDF groups:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF Groups**.
3.
 


 Select an SRDF group, click , and select **Swap Groups**.
4. Select the mirror to refresh.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.



## Setting consistency protection

### Before you begin

To set consistency protection:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Storage Groups > SRDF or DATA PROTECTION > Device Groups > SRDF**.
3.
 



 Select a group, click more , and select **Asynchronous > Set Consistency**.
4. select **Enable** or **Disable**.
5. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration (only applicable for device groups).
6. Click **Advanced Options** to set the advanced [options](#). Select the advanced options and click **OK**.
7. Do one of the following:

- Expand **Add to Job List** and click **Add to Job List Now** 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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Deleting SRDF groups

To delete SRDF groups:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF Groups**.
3. Select the SRDF group and select  .
4. Optional: Click **Advanced Options** and select the **Use Force** check box. This forces the operation.
5. Optional: Click **Advanced Options** and select the **Use SymForce** check box. This forces the operation when the operation would normally be rejected.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing SRDF groups

The SRDF SG list displays a notification if a capacity mismatch exists between R1 and R2 devices. Mismatch can be  $R1 > R2$  or  $R1 < R2$

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > 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**—SRDF group number.

**SRDF Group Label**—SRDF group label, for example, Async, Metro, Witness.

**Remote SRDF group**—Remote RDF group number.

**Type**—Type of group, for example, Dynamic or Witness.

**SRDF Mode**—SRDF modes associated with the SRDF group.

**Online** — Indication if online.

**Transmit Idle** — Time the transmit cycle has been idle.

**Volumes Count**—Number of volumes in the group.

The following controls are available:

 — [Viewing SRDF group details](#) on page 465



**Create SRDF Group**—[Creating SRDF groups](#) on page 457

**Modify**—[Modifying SRDF groups](#) on page 459

**Create Pairs**—[Creating SRDF pairs](#) on page 404



—[Deleting SRDF groups](#) on page 464

**Create SRDF Connection**—[Creating SRDF connections](#) on page 403

**Swap Groups**— [Swapping SRDF groups](#) on page 463

**SRDF/A Pacing Setting**—[Setting SRDF/A pace attributes](#) on page 462

**SRDF/A Setting** —[Setting SRDF/A group attributes](#) on page 461

**SRDF/A DSE Setting**—[Setting SRDF/A DSE attributes](#) on page 461

**Assign Dynamic Cache Partition**—[Assigning dynamic cache partitions](#) on page 945

## Viewing SRDF group details

To view SRDF group details:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION** > **SRDF Groups** to open the **SRDF groups** list view.
3. Select the SRDF group and click  to open its **Details** view.

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 Group Volumes**—SRDF group volumes.

**Director Identity** —Director identifier(s).

**Remote SRDF Group** —Remote group number(s).

**Remote Array ID**—Remote storage system serial ID(s).

**Remote Director Identity**—Remote director identifier(s).

**SRDF Modes**—SRDF Modes. Possible values are: N/A, Adaptive Copy, Synchronous, Asynchronous, Active, and Metro.

**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.

**Transmit Idle Time**—Time the transmit cycle has been idle.

**Transmit Idle Enabled**—Whether SRDF/A Transmit Idle state is active for the SRDF group.

**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/Metro**—SRDF/Metro. Possible values are: Yes, No.

**SRDF/Metro Witness Degraded**—SRDF/Metro Witness Degraded. Possible values are: Yes, No.

**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 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.

**Group Pacing Auto Start**—Indicates if group pacing auto start is enabled/disabled on the SRDF group.

**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.

**SRDF Software Compression**—Indicates if software compression is enabled/disabled on the SRDF group.

**SRDF Single Round Trip** —Indicates if single round trip is enabled/disabled on the SRDF group.

**SRDF 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 Hardware Compression Support** — Indicates if SRDF hardware compression is supported on the storage system.

**Star Mode** — Indicates if SRDF group is in a star configuration.

**SQAR Mode** — Indicates if SRDF group is in a SQAR configuration.

Links are also provided 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 the number next to **SRDF Group Volumes** will open a view listing the 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 402 for additional information.

Before you begin:

SRDF requires Engenuity version 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:

### Procedure

1. Select the storage system.
2. Select **Data Protection > Device Groups**.
3. Click **SRDF**.

The following properties display:

- **Device 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**—[Establishing SRDF pairs](#) on page 421
- **Split**—[Splitting SRDF pairs](#) on page 436
- **Suspend**—[Suspending SRDF pairs](#) on page 436
- **Restore**—[Restoring SRDF pairs](#) on page 433
- **Resume**—[Resuming SRDF links](#) on page 429
- **Failover**—[Failing over](#) on page 422

- **Failback**—[Failing back](#) on page 423
- **Swap**—[Swapping SRDF personalities](#) on page 438
- **Set SRDF/A**—[Setting SRDF/A controls to prevent cache overflow](#) on page 431
- **Set Consistency**—[Setting consistency protection](#) on page 350
- **Move**—[Moving SRDF pairs](#) on page 408
- **Invalidate**—[Invalidating R1/R2 volumes](#) on page 424
- **Ready**—[Making R1/R2 volumes ready](#) on page 425
- **Not Ready**—[Making R1/R2 volumes not ready](#) on page 426
- **R1 Update**—[Updating R1 volumes](#) on page 438
- **RW Enable**—[Read/write enabling R1/R2 volumes](#) on page 428
- **Write Disable**—[Read/write disabling R1/R2 volumes](#) on page 429
- **RW Disable R2s**—[Read/write disabling R2 volumes](#) on page 427
- **Refresh**—[Refreshing R1 or R2 volumes](#) on page 430
- **Set Mode**—[Setting SRDF mode](#) on page 409
- **Delete Pair**—[Deleting SRDF pairs](#) on page 407

## Resuming SRDF links

This procedure explains how to resume I/O traffic on the SRDF links for all remotely mirrored RDF pairs in a group.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > SRDF groups**.
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 , and select **Resume** to open the **Resume** dialog box.

- b. Select the **Use 2nd Hop** option if including the second hop of a cascaded SRDF configuration.
- c. Click **Advanced Options** to set the advanced [SRDF session options](#). Select the advanced options and click **OK**.
- 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

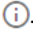
Pair level:

- a. Select a group and click  to open the **SRDF pair list** view.
- b. Select one or more pairs, click **more**, and select **Resume** to open the **Resume** dialog box.
- c. Click **Advanced Options** to set the advanced [SRDF session options](#). Select the advanced options and click **OK**.
- 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing SRDF group volumes

This procedure explains how to view the volumes in an SRDF group:

### Procedure

1. Select the storage system.
2. Select **Data Protection** > **SRDF groups** to open the **SRDF groups** list view.
3. Select the SRDF group and click  to open its **Details** view.
4. Click the number next to **SRDF Group Volumes** to open the **SRDF List Volumes** view.

The following properties display:

**Volumes**—Local volume ID.

**Configuration**—SRDF configuration.

**Remote Symmetrix**—Remote storage system ID.

**Remote SRDF Group**—Remote SRDF group ID.

**Target Volume**—Target volume ID.

**State**—Session state of the pair

**Pair State**—Volume pair state.

**Remote Volume State**—State of the remote volume.

**SRDF Mode**—SRDF copy type.

## SRDF/A control actions

Action	Activate Type	Write Pacing Type	Description
Activate	DSE	N/A	Activates the SRDF/A Delta Set Extension feature, which extends the available

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

## RDFA flags

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

## SRDF group modes

The following values can be set for SRDF groups:

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

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

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

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

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

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

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

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

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

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

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

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

## Understanding RecoverPoint

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

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



## Tagging and untagging volumes for RecoverPoint (storage group level)

### Before you begin

- Volumes that are part of an RDF pair cannot be tagged for RecoverPoint.
- RecoverPoint operations on Unisphere require Enginuity 5876 or higher on the Symmetrix system.
- This feature is not supported on storage systems running HYPERMAX OS 5977.

This procedure explains how to tag (enable) or untag (disable) volumes for RecoverPoint. Enabling volumes makes them accessible to the RecoverPoint appliance.

### Procedure

1. Select the storage system.
2. Select **Storage > Storage Groups**.
3. Do one of the following:
  - 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**.


## Tagging and untagging volumes for RecoverPoint (volume level)

### Before you begin


- Volumes that are part of an RDF pair cannot be tagged for RecoverPoint.
- RecoverPoint operations on Unisphere require Enginuity 5876 or higher on the storage system.
- This feature is not supported on storage systems running HYPERMAX OS 5977.

This procedure explains how to tag (enable) or untag (disable) volumes for RecoverPoint. Enabling volumes makes them accessible to the RecoverPoint appliance.

### Procedure

1. Select the storage system.
2. Select **Storage > Volumes**.
3. In the **All Volumes** panel, expand the type of volume to tag or untag.
4. 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**.
5. Click **OK**.

## Untagging RecoverPoint tagged volumes

### Before you begin

This feature is not supported on storage systems running HYPERMAX OS 5977.

### Procedure

1. Select the storage system
2. Select **DATA PROTECTION > Open Replicator**.
3. Click the **RecoverPoint Volumes** tab.  
Opens the **RecoverPoint Volumes** view.
4. Select a volume and click **Untag**.
5. Click **OK**.


## Viewing RecoverPoint copies

### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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.

This procedure explains how to view the RecoverPoint copies for a particular consistency group.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.  
Opens the **RecoverPoint** list view.
3. Select a RecoverPoint system, click  and click the number next to **Consistency Groups**.  
Opens the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click the number next to **Copies**.  
Opens the **Copies** list view which lists the consistency groups on the selected RecoverPoint system.  
The following properties display:
  - **Copy Name**—Name of copy.
  - **State**—State of the copy. Valid values are *Enabled* or *Suspended*.
  - **Copy Size (GB)**—Size of the copy.
  - **Copy Role**—Current role of the copy. Valid values are *Active* or *Replica*.
  - **RTO (MB)**—Recovery time objective

- **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:



-  —[Viewing RecoverPoint copy details](#) on page 474

## Viewing RecoverPoint copy details

### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.  
Opens the **RecoverPoint** list view.
3. Select a RecoverPoint system, click  and click the number next to **Consistency Groups**.  
Opens the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click the number next to **Copies**.
5. Select a copy and click .

Opens the copy's details view.

The following properties display:

- **Name**—Name of copy.
- **State**—State of the copy. Valid values are *Enabled* or *Suspended*.
- **Role**—Current role of the copy. Valid values are *Active* or *Replica*.
- **Copy Size**—Size of the copy.
- **Journal Size** —Size of the journal, in GB.
- **Journal State** —Indicates the state of the journal. Valid values include *Locked* and *Distributing*.
- **Journal Volume Name** — Indicates the state of the journal. Valid values include *Locked* and *Distributing*.
- **Cluster**—Indicates the state of the journal. Valid values include *Locked* and *Distributing*.
- **RTO (seconds)**—Recovery time objective in seconds
- **Journal Size Limit**—Journal size limit
- **AllowDistribOfLargeSnaps**—Allow distribution of large snapshots
- **AllowSymmWithOneRPA**—Allow storage system with one RPA
- **ActivePrimaryRPA**—Active primary RPA
- **FastForwardBound**—Fast forward bound

- **NumCopySplitters**—Number of copy splitters
- **NumCopyVolumes**—Number of copy volumes
- **NumJournalVolumes**—Number of journal volumes
- **PhoenixDevices**—Phoenix devices
- **TspWritesCleared**—Tsp writes cleared
- **UserSnapshot**—User snapshot
- **Production Copy**—Production copy.
- **Volumes**—Number of associated volumes.
- **Copy Capacity (GB)**—Capacity of the copy, in GB.

## Viewing RecoverPoint sessions

### Procedure

1. Select a storage system.
2. Select **DATA PROTECTION > Open Replicator**.
3. Click the **RecoverPoint Sessions** tab.
4. 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:



— [Viewing RecoverPoint session details](#) on page 475

## Viewing RecoverPoint session details

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Open Replicator**.
3. Click the **RecoverPoint Sessions** tab.
4. Select a session and click .

Opens the session details view.

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 storage groups

### Before you begin

RecoverPoint operations on Unisphere require 5876. 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.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.
3. Select a RecoverPoint system and click  to open the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click the number next to **Copies** to open the **Copies** list view.
5. Select a copy and click  to open the details view.
6. In the properties panel, click the number next to **Storage Groups**.

The following information displays:

**Name**—Name of the storage group.

**Volumes**—Number of volumes in the group.

**Masking views**—Number of associated masking views.

**FAST\_Policy**—FAST policy associated with the RecoverPoint storage group.

**Capacity**—Capacity of the storage group.

**Child SG**—For parent storage groups, this field displays the number of child storage groups; otherwise, this field displays zero.

## Viewing RecoverPoint tagged volumes

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Open Replicator**.
3. Select the **RecoverPoint Volumes** tab 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 controls are available:

 — [Viewing RecoverPoint tagged volume details](#) on page 477

**Untag**—[Tagging and untagging volumes for RecoverPoint \(volume level\)](#) on page 472

## Viewing RecoverPoint tagged volume details

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Open Replicator**.
3. Select the **RecoverPoint Volumes** tab to open the RecoverPoint Volumes list view.
4. Select a volume and click  to open its Details view.

This view allows you to view the volume details.

The following properties display:

**Masking Info**—Number of masking groups.

**Storage Groups**—Number of storage groups.

**FBA Front End Paths**—Number of FBA front end paths.

**RDF Info**—Number of SRDFs.

**Volume 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.

**Encapsulated Device Flag**—Device flag for encapsulated volume. Relevant for external disks only.

**Encapsulated Device Array ID**—Array ID 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 (CYL)**—Volume capacity in cylinders.

**Emulation**—Volume emulation.

**Symmetrix ID**—Storage system on which the volume resides.

**Symmetrix Vol 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 storage system.

**WWN**—World Wide Name of the volume.

**External ID WWN**—External ID 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.

**GCM**—Indicator of GCM.

**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 read miss status.

**System Managed**—The storage system determines the appropriate optimized read miss mode.

## Protecting storage groups using RecoverPoint

### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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 StorageAdmin.
- 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 4.1 is setup and operational. For each cluster in the setup, gatekeepers and repository volumes must be configured in their relevant masking view. uses a default journal masking view naming convention.
- 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.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Storage Groups**.
3. Select the storage group and click **Protect**.
4. On the **Select Technology** page, select **Remote Replication using RecoverPoint**.
5. Click **NEXT**.
6. On the **Configure RecoverPoint** 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.
  - **Data Initiator Group**—Data initiator group.
  - **Journal Thin Pool**—Journal thin pool.

- **Journal Port Group**—Journal port group.
  - **Data Initiator Group**—Journal initiator group.
7. Click **NEXT**.
  8. On the **Add Copies** page, specify the following information:
    - **RecoverPoint Cluster**—RecoverPoint cluster.
    - **Copy Name**—Name of the RecoverPoint copy.
    - **Mode**—Specify whether the mode is Synchronous or Asynchronous.
    - **Array**—Storage system.
    - **Target Storage Group**—Specify whether the RecoverPoint copy targets a new storage group or an existing group.
    - **Copy Storage Group**—Name of storage group to be copied.
    - **Data Thin Pool**—Name of data thin pool.
    - **Data Port Group**—Name of data port group.
    - **Journal Thin Pool**—Name of journal thin pool.
    - **Journal Port Group**—Name of journal port group.
  9. Click **Add Copy**.  
Lists the copy in the **Copy Summary** table.
  10. Click **NEXT**.
  11. On the **FINISH** page, verify your selections. To change any of them, click **BACK**. 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 920 and [Previewing jobs](#) on page 920.
    - Expand **Add to Job List**, then click **Run Now** to perform the operation now.

## Viewing RecoverPoint volumes

### Before you begin

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

To view information on RecoverPoint tagged volumes, refer to [Viewing RecoverPoint tagged volumes](#) on page 476.

To perform this operation, you must be a monitor or higher.




This procedure explains how to view the RecoverPoint volumes for a particular consistency group.

To view RecoverPoint volumes:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems** to open the RecoverPoint Systems view.



3. Select a RecoverPoint system and click .
4. Click the number next to RecoverPoint consistency group.
5. Select a RecoverPoint consistency group and click .
6. Click the number next to Replication Sets.
7. Select the replication set and click .
8. Click the number next to Volumes.

The following properties display:

**Volume Name**—Name of the 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 ID**—Array ID.

**Vendor**—Vendor of the volume.


**Product Name**—Storage product installed.

## Viewing RecoverPoint clusters

### Before you begin

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

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.  
Opens the **RecoverPoint Systems** list view.
3. Select a **RecoverPoint** system.
4. Click  and click the number next to **Clusters**.



Opens the **RecoverPoint Clusters** table view.

The following information displays:

- **Cluster Name**—Name of the cluster.
- **RecoverPoint Appliances**—Number of RecoverPoint appliances.
- **IPv4 Address**—IP address, in IPv4 format. If an IPv6 address is used, this column has the value "N/A".
- **IPv6 Address**—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

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.  
Opens the **RecoverPoint Systems** list view.
3. Click  and click the number next to **Clusters**.  
Opens the cluster list view.
4. Select a cluster and click .  
Opens the cluster details view.  
The following properties display:
  - **Cluster Name**—Volume name.
  - **IPv4 Address**—IP address, in IPv4 format. If an IPv6 address is used, this column has the value "N/A".
  - **IPv6 Address**—IP address, in IPv6 format. If an IPv4 address is used, this column has the value "N/A".
  - **RecoverPoint Appliances**—Number of RecoverPoint appliances.
  - **RecoverPoint Splitters**—Number of RecoverPoint splitters.
  - **Software Serial ID**—Serial ID of the software.
  - **RPA Type**—RecoverPoint appliance type.
  - **Timezone** —Time zone.
  - **Maintenance Mode**—Maintenance mode in use.
  - **Internal Cluster Name**—Internal name of the cluster.

## Viewing RecoverPoint splitters



### Before you begin

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

This procedure explains how to view RecoverPoint splitters.

To view RecoverPoint splitters:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.  
Opens the **RecoverPoint Systems** list view.
3. Select a RecoverPoint system.
4. Click  and click the number next to **Clusters**.  
Opens the cluster list view.
5. Select a cluster and click .

Opens the cluster details view.

6. Click the number next to **Splitters** to open the **RecoverPoint Splitters** list view.

The following information displays:

**Name**—Name of the splitter.

**Array ID**—Array ID of the splitter.

**Array Type**—Array type of the splitter.

**Status**—Status of the splitter.

**Attached RPA Cluster**—Number of attached clusters.

## Viewing RecoverPoint appliances


### Before you begin

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


### Procedure


1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.

Opens the **RecoverPoint Systems** list view.

3. Select a RecoverPoint system.
4. Click  and click the number next to **Clusters**.

Opens the cluster list view.

5. Select a cluster and click .
- Opens the cluster details view.

6. Click  and click the number next to **RecoverPoint Appliances**.

Opens the RecoverPoint Appliances view and displays the following information:

- **Name**—Name of the RecoverPoint appliance.
- **Status**—Status of the RecoverPoint appliance.
- **WAN (IP)**— Wide Area Network (WAN) IP address
- **Management IPv4** — IP address, in IPv4 format.
- **Local Fibre Connectivity**—Local RPA Fibre Connectivity
- **Remote Fibre Connectivity**—Remote RPA Fibre Connectivity

## RecoverPoint systems

### Manage RecoverPoint discovery

To discover a RecoverPoint system, see [Discovering RecoverPoint Systems](#) on page 484.

To update RecoverPoint discovery information, see [Updating RecoverPoint discovery information](#) on page 484.

## Discovering RecoverPoint Systems

### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- This operation requires StorageAdmin privileges.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.
3. Click **Create**.
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.
  - **Port**—System port number.
  - **System Username**—System username.
  - **System Password**—System password.
  - **Confirm System Password**—Re-enter system password.
5. Click **OK**.

## Deleting RecoverPoint systems

### Before you begin

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

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.
3. Select a RecoverPoint system and click **Delete RecoverPoint System**.
4. Click **OK**.

## Updating RecoverPoint discovery information

### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 on the storage system. RecoverPoint operations are not supported on storage systems running HYPERMAX OS 5977 or higher.
- This operation requires StorageAdmin privileges.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.
3. Select a RecoverPoint system.

4. Click **Update Discovery Information**.
5. Type the following information.
  - **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

### Before you begin

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

This procedure explains how to view previously discovered RecoverPoint systems.

To view RecoverPoint systems:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems** to open the RecoverPoint Systems view.

The following properties display:

**System Name**—Name of the system.

**IPv4 Address**—IP address of the system.

**Port**—Port 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.

The following controls are available:

 —[Viewing RecoverPoint system details](#) on page 485

**Create**—[Discovering RecoverPoint Systems](#) on page 484

**Update Discovery Information**—[Updating RecoverPoint discovery information](#) on page 484

**Delete RecoverPoint System**—[Deleting RecoverPoint systems](#) on page 484


## Viewing RecoverPoint system details

### Before you begin

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

To view RecoverPoint system details:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems** to open the RecoverPoint Systems view.
3. Select the system and click  to open its Details view.

The following properties display:

**RecoverPoint Systems**—Name of the system.

**Consistency Groups**—Number of consistency groups associated with the system.

**Clusters**—Number of RPA clusters in the system.

**Critical Alerts Count**—Number of critical alerts.

**OK Alerts Count**—Number of OK alerts.

**Warning Alerts Count**—Number of warning alerts.

**Critical Events Count**—Number of critical events.

**Warning Events Count**—Number of warning events.

**Events Error Count**—Number of events errors.

## RecoverPoint consistency groups

### Viewing RecoverPoint consistency groups

#### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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.

This procedure explains how to view the consistency groups used to protect the RecoverPoint volumes.

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint**.

Opens the **RecoverPoint** list view.

3. Select a RecoverPoint system, click  and click the number next to **Consistency Groups**.

Opens the **Consistency Group** list view which lists the consistency groups on the selected RecoverPoint system.

The following properties display:

- **Consistency Group**—Consistency group name.
- **Group Enabled**—Consistency group state.
- **Link States**—Lists the states of associated links.
- **Source Capacity (GB)** —Source capacity in GB.
- **Primary RPA**—Primary RecoverPoint appliance number.

- **Production Copy**—Name of the production copy.

The following controls are available:

- ⓘ —[Viewing RecoverPoint consistency group details](#) on page 487
- **Copies**—[Viewing RecoverPoint copies](#) on page 473
- **Replication Sets**—[Viewing RecoverPoint replication sets](#) on page 488
- **Active Links**—[Viewing RecoverPoint links](#) on page 489

## Viewing RecoverPoint consistency group details

### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint**.  
Opens the **RecoverPoint** list view.
3. Select a RecoverPoint system, click ⓘ and click the number next to **Consistency Groups**.
4. Select a consistency group and click ⓘ to view the properties of that Consistency Group.

Displays the properties of the Consistency Group.

The following properties display:

- **Group State**—State of the group.
- **Group Setting**—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**—Distributed group.
- **Managed by RecoverPoint** —Indicates if the consistency group is managed by RecoverPoint.
- **Read Only Replica Volumes** —Read-only replica volumes.

## RecoverPoint replication sets

### Viewing RecoverPoint replication sets

#### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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.

This procedure explains how to view the RecoverPoint replication sets for a particular consistency group.

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint**.

Opens the **RecoverPoint** list view.

3. Select a RecoverPoint system, click  and click the number next to **Consistency Groups**.

Opens the **Consistency Group** list view.

4. Select a RecoverPoint consistency group and click the number next to **Replication Sets**.

Opens the **Replication Sets** list view, which lists 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 control is available:

-  — [Viewing RecoverPoint replication set details](#) on page 488

### Viewing RecoverPoint replication set details

#### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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.



This procedure explains how to view the details of a RecoverPoint replication set.

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint**.

Opens the **RecoverPoint** list view.



3. Select a RecoverPoint system, click  and click the number next to **Consistency Groups**.  
Opens the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click the number next to **Replication Sets**.
5. Select a replication set, and click .  
The following properties display:
  - **Name**—Name of the replication set.
  - **Volumes**—Number of associated volumes.
  - **Volume Name**—Name of associated volume.
  - **Production Volume Capacity (GB)**—Production volume capacity, in GB .
  - **Capacity (GB)**—Source capacity, in GB.

## RecoverPoint links


### Viewing RecoverPoint links

#### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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.

This procedure explains how to view the RecoverPoint links for a particular consistency group.

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint Systems**.  
Opens the **RecoverPoint Systems** list view.
3. Select a RecoverPoint system, click  and click the number next to **Consistency Groups**.  
Opens the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click the number next to **Links**.  
Opens the **Links** list view, which lists the links associated with the selected consistency group.  
The following properties display:
  - **Name**—Name of the RecoverPoint link.
  - **Transfer Enabled** —Indicates if the transfer state is enabled for this RecoverPoint link.
  - **Link State**—Current role of the copy. Valid values are *Active* or *Replica*.
  - **Local Link**—Indicates if the link state is active or paused.
  - **Protection Mode**—Protection Mode.
  - **RPO (seconds)**—RPO.

The following control is available:



-  — [Viewing RecoverPoint link details](#) on page 490

## Viewing RecoverPoint link details

### Before you begin

- RecoverPoint operations on Unisphere require Enginuity 5876 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.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > RecoverPoint**.  
Opens the **RecoverPoint** list view.
3. Select a RecoverPoint system, click  and click the number next to **Consistency Groups**.  
Opens the **Consistency Group** list view.
4. Select a RecoverPoint consistency group and click the number next to **Links**.
5. Select a link and click .

Opens the link's details view.

The following properties display:

- **Name**—Name of the RecoverPoint link.
- **Transfer State**—Indicates if the transfer state is enabled for this RecoverPoint link.
- **Link State**—Indicates if the link state is active or paused.
- **Local**—Indicates if the link is local.
- **RPO (seconds)**—Recovery point objective in seconds.
- **First Copy**—First copy.
- **Second Copy**—Second copy.
- **Protection Mode**—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**—Bandwidth limit.
- **Deduplication**—Specifies if deduplication is enabled.
- **Snapshot Granularity**—Snapshot granularity.

## Creating Open Replicator copy sessions

### Before you begin

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

- Push
- Differential
- Precopy

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

### Procedure

1. Select the storage system.
2. Select **Data Protection > Open Replicator > Open Replicator SAN View**.
3. Filter the items displayed in the **Filtered LUNs** panel by selecting items within the **Control Ports**, **Remote Ports**, and **Remote Volumes** panels.
4. Select one or more volumes within the **Filtered LUNs** panel and click **Create Copy Session**.
5. Click **Create Copy Session**.
6. Select a **Copy Direction** and **Copy Operation**.
7. 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.

8. Select a remote volume and target volume, then click **Add Pair**.

If the pair is valid, it is added to the **Volume Pairs** list.

9. Click **Remove Pair** to edit the **Volume Pairs** list.
10. Click **Next**.
11. Enter **Session Name**.
12. 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.

13. Select the [Open Replicator session options](#) and click **Next**.
14. View session **Summary** and click **Finish** to create session or click **Back** to edit session options.

## Activating Open Replicator session

### Before you begin

The copy session must be in a created or recreated state before you can activate it.

### Procedure

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 the **Activate Session** dialog box.
4. Select a copy option.  
Refer to [Open Replicator session options](#) for session copy and control options.
5. Click **OK**.

## Recreating Open Replicator sessions

### Before you begin

Recreating operations are not supported when the ORS control volume is on a storage system running HYPERMAX OS 5977.

### Procedure

1. Select the storage system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions**.
3. Select a session and click **Recreate**.
4. Optional: Select the **PreCopy** or **Force** checkbox or both checkboxes.
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.

### Procedure

1. Select the storage system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions**.
3. Select a session and click **Restore**.
4. Select any number of the available options. Refer to [Open Replicator session options](#) 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.

### Procedure

1. Select the storage system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions**.
3. Select a session and click **Rename**.

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.

### Procedure

1. Select the storage system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions**.
3. Select a session and click **Remove**, and click **OK**.

An error message is displayed if the session is in a state that does not allow the session to be removed.

## Setting Open Replicator session background copy mode

### 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.

This procedure sets the session background copy mode for an ORS session that has already been created.

### Procedure

1. Select the storage system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions**.
3. Select a session and click **Set Mode**.
4. Select the background copy mode. Refer to [Open Replicator session options](#) 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.

### Procedure

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 the **Set Donor Update Off** dialog box.
4. Select the [Open Replicator session options](#).
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.

**Procedure**

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](#) for session control options.

4. Click **OK**.

**Setting Open Replicator session pace****Before you begin**

This feature is not supported on storage systems running HYPERMAX OS 5977 or higher.

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.

**Procedure**

1. Select the storage system.
2. Select **Data Protection > Open Replicator > Open Replicator Sessions**.
3. Select a session and click **Set Pace**.
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.

**Procedure**

1. Select a storage system.
2. Select **System > 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****Procedure**

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](#) for session control options.
5. Click **OK**.

## Viewing Open Replicator sessions

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Open Replicator** and click **Open Replicator Sessions**.

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.

The following controls are available:

- ⓘ —[Viewing Open Replicator session details](#) on page 496
- **Activate**—[Activating Open Replicator session](#) on page 491
- **Terminate**—[Terminating Open Replicator sessions](#) on page 494
- **Front End Zero Off**—[Setting Open Replicator session donor update off](#) on page 493
- **Donor Update Off**—[Setting Open Replicator session donor update off](#) on page 493
- **Rename**—[Renaming Open Replicator sessions](#) on page 492 This option is not available for systems running HYPERMAX OS 5977 or higher.
- **Remove**—[Removing Open Replicator sessions](#) on page 493 This option is not available for systems running HYPERMAX OS 5977 or higher.
- **Restore**—[Restoring Open Replicator sessions](#) on page 492 This option is not available for systems running HYPERMAX OS 5977 or higher.
- **Recreate**—[Recreating Open Replicator sessions](#) on page 492 This option is not available for systems running HYPERMAX OS 5977 or higher.
- **Set Mode**—[Setting Open Replicator session background copy mode](#) on page 493
- **Set Pace**—[Setting Open Replicator session pace](#) on page 494 This option is not available for systems running HYPERMAX OS 5977 or higher.

## Viewing Open Replicator session details

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Open Replicator** and click **Open Replicator Sessions**.
3. Select a session and click  to open the session details view.

Depending on the configured system, some or all of the following properties display:

- **Session**—ORS session name.
- **Control Volume**—Control volume name.
- **Remote Volume**—Remote volume name.
- **Remote Volume Specification**— Remote volume specification. (Not applicable for storage systems running HYPERMAX OS 5977 or higher.)
- **Status**—Session status.
- **Percent Complete**— Percent tracks copied. (Not applicable for storage systems running HYPERMAX OS 5977 or higher.)
- **Copy Pace**—Copy Pace value (0 - slowest to 9 - fastest, default is 5). (Not applicable for storage systems running HYPERMAX OS 5977 or higher.)
- **Protected Tracks**—Number of protected tracks.
- **Modified Tracks**—Number of modified tracks. (Not applicable for storage systems running HYPERMAX OS 5977 or higher.)
- **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.
- **RecoverPoint Session**—Indicates if session is a RecoverPoint session. (Not applicable for storage systems running HYPERMAX OS 5977 or higher.)
- **Standard ORS Session**—Indicates if session is a standard session. (Not applicable for storage systems running HYPERMAX OS 5977 or higher.)
- **Front-End Zero**—Indicates if front-end zero detection is enabled.

## Viewing Open Replicator SAN View

### Procedure

1. Select the storage system.
2. Select **Data Protection > Open Replicator > Open Replicator SAN View**.

Use this view to view select remote volumes in the **Filtered LUNs** panel to use for Open Replicator copy sessions. The list of volumes can be filtered further by



selecting items within the **Control Ports**, **Remote Ports**, and **Remote Volumes** panels.

The following controls are available:

- **Create Copy Session**—[Creating Open Replicator copy sessions](#) on page 490
- **Rescan**—Causes a rescan operation to be performed.

## 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.
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.

Session Option	Used with Command	Description
Force	Terminate Restore	Select the <b>Force</b> option if the copy session is in progress. This will allow the session to continue to copy in its current mode without donor update.
	Donor Update Off	Select the <b>Force</b> option if the copy session is in progress. This will allow the session to continue to copy in its current mode without donor update.
Force Copy	Activate	Overrides any volume restrictions and allows a data copy.  For a push operation, remote capacity must be equal to or larger than the control volume extents and vice versa for a pull operation. The exception to this is when you have pushed data to a remote volume that is larger than the control volume, and you want to pull the data back, you can use the <b>Force_Copy</b> option.
Front-End Zero Detection	Create	Enables front end zero detection for thin control volumes in the session. Front end zero detection looks for incoming zero patterns from the remote volume, and instead of writing the incoming data of all zeros to the thin control volume, the group on the thin volume is de-allocated.
Hot	Create	Hot copying allows the control device to be read/write online to the host while the copy operation is in progress. All directors that have the local devices mapped are required to participate in the session. A hot copy session cannot be created unless all directors can discover the remote device.
Nocopy	Activate	Temporarily stops the background copying for a session by changing the state to CopyOnAccess or CopyOnWrite from CopyInProgress.
Pull	Create	A pull operation copies data to the control device from the remote device.

Session Option	Used with Command	Description
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. R = RecoverPoint session. S = Standard ORS session.
Z Front-end zero detection	X =Enabled . = Disabled
*	Failed session can be reactivated.

## Understanding non-disruptive migration (NDM)

Non-disruptive migration (NDM) allows you to migrate storage group (application) data in a non-disruptive manner with no downtime from source arrays running Enginuity OS 5876 Q3 2016 or higher to target arrays running HYPERMAX OS 5977 Q3 2016 or higher.

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

NDM applies to open systems/FBA devices only.

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

From Unisphere 8.4 onwards, an NDM session can be created on a storage group containing session target volumes (R2s) where the SRDF mode is synchronous. The target volumes of an NDM session may also have a SRDF/Synchronous session added after the NDM session is in the cutover sync state.

The following NDM tasks can be performed from Unisphere.

[Setting up a migration environment](#)

Optional: [Removing a migration environment](#)

Optional: [Preparing a NDM session](#)

[Creating a NDM session](#)

[Viewing NDM sessions](#)

[Viewing NDM session details](#)

[Cutting over a NDM session](#)

Optional: [Stop synchronizing data after NDM cutover](#)

Optional: [Start synchronizing data after NDM cutover](#)

[Committing a NDM session](#)

Optional: [Cancelling a NDM session](#)

Optional: [Recovering a NDM session](#)

[Viewing migration environments](#)

[Adding a migration environment](#)

[Removing a migration environment](#)

## Preparing a non-disruptive migration (NDM) session

Non-disruptive migration of storage groups using SRDF is supported between a source storage system running Enginuity 5876 Q3 2016 or higher and a target storage system running HYPERMAX OS 5977 Q3 2016 or higher.

See [Understanding non-disruptive migration \(NDM\)](#) on page 500 for additional information.

There are two paths through the migration creation wizard. The default flow is for creating a migration session between two arrays (see [Creating a non-disruptive migration \(NDM\) session](#) on page 502). The secondary flow allows the user to prepare for a data migration with recommendations on the ports to be used for an optimal candidate migration result. When the prepare path is run (this is an option that can be run before the create path), you have the option to save your preparation to a Migration report containing zoning information. You need to implement the zoning

before running the Create scenario in anticipation of the migration creation. If the plan is changed after running the prepare, these port groups need to be renamed or removed.

If the user chooses the prepare path first, the same Symmetrix and SRP must be selected when running the second path for creating the actual migration session.

Before you begin:

To perform this procedure you must be an Administrator or Storage Admin.

The data migration environment exists between two candidate arrays.

The selected storage group is a masked candidate storage group.

The selected storage group does not contain only gatekeepers.


The local array must have online RDF ports.

Unisphere is registered for performance data processing on the source and target arrays. 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 Unisphere GUI and charts.

To prepare a migration session:

#### Procedure

1. Select a storage system running Enginuity 5876 Q3 2016 or higher.
2. Select **STORAGE > Storage Groups**.
3. Select a storage group.
- 4.

Click  and click **Migrate**.

5. Select the target storage system.
6. Select the target Storage Resource Pool (SRP).

Not specifying an SRP is allowed for data migration creation.

7. Select a port group.
8. Click **NEXT**.
9. Do the following:

- Select **Prepare Data Migration**

If the source or target array is remote to this instance of Unisphere, performance data processing is not registered on the target array, or there has not been sufficient time (at least two intervals (by default, 5 minutes each)) to gather performance data, an error popup informs you of this and the **NEXT** button is disabled.

If any source port groups do not already exist on the target array, a panel is displayed allowing the user to select ports for any port group(s) to be created.

All port group(s) involved in this migration are displayed. Any port group(s) that need to be created on the target array are at the top and any that already exist are at the bottom. Any existing port group(s) have the text "Already configured" in the title.

Any port group to be created displays a selectable list of ports. This list of ports includes all available ports on the target array, but to avoid overlap,

port(s) already in use by any existing target array port group(s) are filtered out of the list.

The port table within the panel contains the following columns:

- **Port**—The port identifier of a target array port in Dir:Port format with a checkbox for selection.
- **Utilization**— a bar indicating a utilization score for the port. A lower score indicates lower utilization. This is the default sort column for the list.
- **Initiators**— a number indicating how many initiators, from the list of all initiators in the corresponding source Masking View associated with the source Storage Group, are present in the Login History Table for the port on the target array.

Ports are selected by default based on the Utilization value. The number of default selected ports is equal to the number of ports in the source port group or the number of ports still available in the original list. You are able to override these selections, but you must select at least one port.

- Click **NEXT**.
- On the **Summary** page, review the details. The summary includes information any port group(s) and ports that you selected. There is also a suitability score for the entire migration request indicating the expected impact of the migrated application on the target array's front end ports. A message, indicating whether or not the selected front end ports have sufficient performance capacity for the incoming load, is displayed. Do one of the following:
  - Optional: Click **Save Migration report** to save the report to your chosen location.  
You need to implement zoning based on the information in the Migration report. You need to implement zoning before running the Create scenario as well as creating the required port groups on the target array in anticipation of the migration creation. If the plan is changed after running the prepare, these port groups need to be renamed or removed.
  - Click **Finish** to perform the port group(s) creation (if any) on the target array depending on your selections.

---

#### Note

Clicking **Finish** does not create the migration session.

---

## Creating a non-disruptive migration (NDM) session

Non-disruptive migration of storage groups using SRDF is supported between a source storage system running Enginuity 5876 Q3 2016 or higher and a target storage system running HYPERMAX OS 5977 Q3 2016 or higher.

When migrating a storage system from HYPERMAX OS 5977 to PowerMaxOS 5978, a create with precopy option is supported by Unisphere. The precopy option allows storage to be provisioned on the target array without making the devices' host visible. This allows the application to continue running on the source array while data is being copied to the target.

See [Understanding non-disruptive migration \(NDM\)](#) on page 500 for additional information.

There are two paths through the migration creation wizard. The default flow is for creating a migration session between two arrays. The secondary flow will allow the user to prepare for a data migration (see [Preparing a non-disruptive migration \(NDM\) session](#) on page 500).

Before you begin:

To perform this procedure you must be an Administrator or Storage Admin.

The data migration environment exists between two candidate arrays.

The selected storage group is a masked candidate storage group.

The selected storage group does not contain only gatekeepers.

The Initiators in the Storage Groups Masking Views are visible to the target array running HYPERMAX OS 5977 or higher.

The local array must have online SRDF ports.

You are allowed to select a port group name on the target array to use as part of the migrated Masking View. This port group must exist on the target array.


When migrating a storage system from HYPERMAX OS 5977 to PowerMaxOS 5978, a create with precopy option is supported by Unisphere. The precopy option allows storage to be provisioned on the target array without making the devices' host visible. This allows the application to continue running on the source array while data is being copied to the target.

To create a migration session:

#### Procedure

1. Select a storage system running Enginuity OS 5876 Q3 2016 or higher.
2. Select **STORAGE > Storage Groups**.
3. Select a storage group.
- 4.



Click  and click **Migrate**.

5. Select the target storage system.
6. Select the target Storage Resource Pool (SRP).

The default SRP is selected on the SRP combo if it can be calculated. Not specifying an SRP is allowed for data migration creation.

7. Select a port group.
8. Click **NEXT**.
9. Do the following:
  - Select **Create Data Migration**
  - Optional: Uncheck the **Compression** check box to turn off Compression. Compression is only allowed on All Flash systems running the HYPERMAX OS 5977 Q3 2016 Service Release or higher.
  - Optional: Click **Precopy**
  - Click **NEXT**.
  - On the **Summary** page, review the details. The summary includes information on any Masking View(s) that would be created by this migration and any Port group(s) and Host/Host Group(s) that you selected.
  - 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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now. Review the contents of the feedback dialog. After successful migration, a dialog is displayed. Select **Go to Migrations list view**, **No Further action at this time** or **Close**.

### Results

If the host can scan the new paths on its own, the migration moves to the **CutoverReady** state. If a user rescan is needed, the migration state moves to the **Created** state.

## Viewing the non-disruptive migration (NDM) sessions list

This procedure explains how to view the list of the non-disruptive migration (NDM) sessions.

See [Understanding non-disruptive migration \(NDM\)](#) on page 500 for additional information.

To view the migration sessions list:


### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Select the **Storage Groups** tab.

The following properties display:

- **Storage Group**—Name of the storage group.
- **State**—Migration state. An icon representing the state is also displayed. Failed states are represented as red, in progress states are represented using the refresh icon and states after successful completion of actions are green.
- **Source**—Source storage system.
- **Target**—Target storage system.

The following controls are available:

-  — [Viewing migration details](#) on page 504
- **Cutover**—[Cutting over a migration session](#) on page 506
- **Commit**—[Committing a migration session](#) on page 507
- **Ready Target**—[Readying the migration target](#) on page 506
- **Recover**—[Recovering a migration session](#) on page 509
- **Sync**—[Synchronizing data after non-disruptive migration \(NDM\) cutover](#) on page 507
- **Cancel Migration**—[Cancelling a migration session](#) on page 508

## Viewing migration details


This procedure explains how to view the migration details for a specific data migration.



See [Understanding non-disruptive migration \(NDM\)](#) on page 500 for additional information.

To view the migration details:

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Select the **Storage Groups** tab.
4. Select a storage group and click  to view the **Migrations** details view.

The following items are displayed:

- **Storage Group**—Name of the storage group.
- **State**—Migration state.
- **Source**—Source storage system.
- **Target**—Target storage system.
- **Capacity (GB)**—Capacity of the storage group in GB.
- **Synched Capacity (GB)**—Synchronized capacity of the storage group in GB.
- A storage group table displaying the source status and target status for each storage group associated with the migration.
- A masking view table displaying the source status and target status for each masking view associated with the migration.
- A Port Group table. Selecting a row in the masking view table populates the Port Group table. The table displays the source status and target status for the selected masking view.
- A Host/Host Group table. Selecting a row in the masking view table populates the Host/Host Group table. The table displays the source status and target status for the selected masking view..

Select an item in the Storage Group table to view the following volume information:

- **Source Volume**—Identity of source volume.
- **Source Status**—Status of the source volume..
- **Target Volume**—Identity of target volume.
- **Target Status**—Status of the target volume.

Select an item in the Port Group table to view the following port information:

- **Symmetrix**—Storage system ID.
- **Port Name**—Identity of port.
- **Status**—Status.

Select an item in the Host/Host Group table to view the following Initiators information:

- **Initiator**—Identity of initiator.
- **Source Status**—Source status.

- **Target Status**—Target status.

## Readying the migration target

### Before you begin

The migration must be in the precopy state.

To perform this procedure you must be a Storage Admin.

This operation is used on migrations that are in the Precopy state. When migrating a storage system from HYPERMAX OS 5977 to PowerMaxOS 5978, a create with precopy option is supported by Unisphere. The precopy option allows storage to be provisioned on the target array without making the devices' host visible. This allows the application to continue running on the source array while data is being copied to the target. The Ready Target operation results in the target devices becoming visible to the host and configures the data migration to allow simultaneous access to both the source and target devices.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Click the **Storage Groups** tab.
- 4.

Select the storage group, click , and click **Ready Target**.

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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

### Results

If the operation is successful, a success message appears indicating the Ready Target operation was successful and that a host discovery needs to be performed. The state at this stage is Migrating. Once the host discovery has been performed and all data synchronized between the source and target arrays, the migration state changes to Synchronized.

If the command was unsuccessful an error message will appear detailing the reason for the command failure. If the Ready Target operation has run to completion with a failed status the migration has a status of 'Ready Target Failed'.

## Cutting over a migration session

The cutover operation results in the storage array running HYPERMAX OS 5977 Q3 2016 or higher becoming the active array.

See [Understanding non-disruptive migration \(NDM\)](#) on page 500 for additional information.

Before you begin:

To perform this procedure you must be an Administrator or Storage Admin.

The state of the migration session is CutoverReady.

To cutover a migration session:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Select the **Storage Groups** tab.
4. Select a storage group and click **Cutover**.
5. (Optional) Under the **Advanced Options** link you can select **Force** and/or **SymForce**.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Synchronizing data after non-disruptive migration (NDM) cutover

After cutover, the sync operation allows the synchronization of the data between an array running HYPERMAX OS 5977 Q3 2016 or higher and an array running Enginuity 5876 Q3 2016 or higher.

See [Understanding non-disruptive migration \(NDM\)](#) on page 500 for additional information.

Before you begin:


To perform this procedure you must be an Administrator or Storage Admin.


The state of the migration session is CutoverSync, CutoverSyncing or CutoverNoSync.

To start or stop the synchronization of a migration session:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Select the **Storage Groups** tab.
4.
 



 Select the storage group, click , and click **Sync**.
5. Click **Stop** to stop a synchronization or **Start** to start a synchronization.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Committing a migration session

This procedure explains how to commit a data migration session.

See [Understanding non-disruptive migration \(NDM\)](#) on page 500 for additional information.

Before you begin:

To perform this procedure you must be an Administrator or Storage Admin.

The state of the migration session is CutoverSync or CutoverNoSync.

To commit a migration session:

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Select the **Storage Groups** tab.
4. Select a storage group and click **Commit**.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Cancelling a migration session

This procedure explains how to cancel a data migration session (from the source or target array).

See [Understanding non-disruptive migration \(NDM\)](#) on page 500 for additional information.

Before you begin:

To perform this procedure you must be an Administrator or Storage Admin.

There are two cancel options. One of them is chosen by Unisphere depending on the migration session state.


The "Cancel without the revert flag" set operation requires the state of the migration session to be: Created, CreateFailed or CutoverReady.

The "Cancel with the revert flag" set operation requires the state of the migration session to be: Migrating or CutoverSync.

To cancel a migration session:

#### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Select the **Storage Groups** tab.
- 4.

Select the storage group, click , and click **Cancel**.

A checkbox is displayed when Revert is being used and a dash is displayed when Revert is not being used.

5. (Optional) Select the **SymForce** check box.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Recovering a migration session

This procedure explains how to recover a data migration session from a failed state.

See [Understanding non-disruptive migration \(NDM\)](#) on page 500 for additional information.

Before you begin:


To perform this procedure you must be an Administrator or Storage Admin.

The state of the migration session is CreateFailed, CutoverFailed, CancelFailed, MigrFailed, or RevertFailed.

To recover a migration session:

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Select the **Storage Groups** tab.
- 4.

Select the storage group, click , and click **Recover**.

5. (Optional) Select **Force** in the **Advanced Options** section.

The force flag is meant for cases where the state of the migration session is CreateInProg, CancelInProg, CutoverInProg, RevertInProg, or CommitInProg.

6. (Optional) Select **SymForce** in the **Advanced Options** section.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing migration environments

This procedure explains how to view existing migration environments. You can also use this view to add a new migration environment or delete an existing one:

- [Adding a migration environment](#)
- [Removing a migration environment](#)

### Procedure


1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Select the **Environments** tab

The following properties display:

- **Source**—Source storage system.
- **Target**—Target storage system.
- **State**—Migration state. An icon representing the state is also displayed. Failed states are represented as red, in progress states are represented using the refresh icon and states after successful completion of actions are green.

- **In Use**—Provides an icon that indicates whether any objects involved in the migration session are in use

The following controls are available:

- **Create**—[Setting up a migration environment](#) on page 510
-  —[Removing a migration environment](#) on page 510

## Setting up a migration environment

### Before you begin

To perform this procedure you must be a Storage Admin.

This procedure explains how to add a new migration environment.

### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Select the **Environments** tab and click **Create**.
4. Select the remote storage system for the environment creation.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.


## Removing a migration environment

### Before you begin

To perform this procedure you must be a Storage Admin.

This procedure explains how to remove a migration environment.



### Procedure

1. Select the storage system.
2. Select **DATA PROTECTION > Migrations**
3. Click the **Environment** tab, select a migration environment in the list view and click  .
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing the authorized users and groups details

To view the authorized users and groups list:

**Procedure**

1. Select  to open the **Settings** panel.
2. Select **Users and Groups > Authorized Users & Groups**.
3. Select your required storage system ID from the drop-down list.
4. To see more information on a user, select the user and on the right-hand side of the row, click the  icon.
5. View the following information in the information dialog: name, authority, domain, storage system identity, roles, and component name.

## Expanding remote volumes

**Before you begin**

- Requires HYPERMAX OS 5977 or later (HYPERMAX OS 5977.1125.1125 or later for CKD volumes). Consider consulting with your operating system vendor or cluster vendor for support of online LUN expansion

**Procedure**



1. Selects a group (one only) to expand remote volume or do not select a group to only expand local volume.
2. Click **APPLY**.

## Setting a device identity

**Before you begin**

This operation allows the conversion of device IDs between Compatibility IDs and Mobility IDs.

**Procedure**

1. Select the storage system.
2. Select **STORAGE > Volumes**
3.   
Select a FBA volume, click , and click **Set Volumes > Mobility ID**.
4. Select **Mobility ID** if you wish to change the device identity type from Compatibility ID (default) to Mobility ID.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

**Results**

If the operation is successful, a success message appears indicating the Ready Target operation was successful and that a host discovery needs to be performed. The state at this stage is Migrating. Once the host discovery has been performed and all data synchronized between the source and target arrays, the migration state changes to Synchronized.

If the command was unsuccessful an error message will appear detailing the reason for the command failure. If the Ready Target operation has run to completion with a failed status the migration has a status of 'Ready Target Failed'.

## Editing storage group volume details

To edit storage group details for a storage system running Hypermax OS 5977 or higher:

### Procedure

1. Click the **Volume Config** tab.
2. 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 modifying storage groups with new volumes. Note that when modifying 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.

3. Optional: Click **Enable Compression** checkbox.
4. Optional: Click the **Enable Mobility ID** checkbox to assign Mobility IDs to the volumes in the storage group. If you leave the checkbox unchecked, Compatibility IDs will be assigned to the volumes instead.
5. Optional: Click **Allocate capacity for each volume** checkbox.
6. Optional: Click **Persist preallocated capacity through reclaim or copy** checkbox.
7. Click the **Volume Size** tab.
8. Enter a volume size, capacity and capacity unit.
9. Optional: Add one or more volume sizes by hovering over the area to the right of the volume capacity and selecting .
10. Optional: Click  to remove a volume size.
11. Click **APPLY**.

The Storage Group page in the wizard displays **Mixed Capacities** for the row. Click **Mixed Capacities** to reopen this dialog.



## Editing storage group details

To edit storage group details for a storage system running Enginuity 5876:

### Procedure

1. 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.
- 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. (applicable for regular volumes only)

2. Click **OK**.

## Replication state severities

The following are the replication state severities, organized by replication technology:

Severity	SRDF State	TimeFinder State
Normal	Synchronized	Copied
	SyncInProg	Created
	Split	Restored
	Failed Over	TerminateInProgress
	R1 UpdInProg	Recreated
	R1 Update	CopyInProgress
	Consistent	CopyOnAccess
		CreateInProgress
		CopyOnWrite
		RestoreInProgress
		PreCopy

Severity	SRDF State	TimeFinder State
		Split SyncInProgress SplitInProgress SplitNoIncremental EstablishInProgress Linked LinkCopied NeverEstablished Synchronized SplitBeforeSync SplitBeforeRestore Established LinkCopyInProgress
Warning	Suspended TransIdle	Warning
Error	Partitioned Mixed N/A Invalid	N/A Invalid Failed

## Managing space reclamation

### Before you begin






- You can only reclaim thin pool capacity from bound thin volumes.
- Thin pool reclamation for individual thin volumes requires Enginuity 5876 or HYPERMAX OS 5977 or higher.
- Thin pool reclamation for thin volumes in a storage group requires Enginuity 5876 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.

The following describes how to start and stop the space reclamation option, freeing unused pool capacity and reclaiming unwritten tracks from thin volumes. In addition, you can also perform this operation from the following views:

- Storage Groups (HYPERMAX OS 5977 or higher): **STORAGE > > Storage Groups**
- Storage Groups (Enginuity 5876): **STORAGE > Storage Groups**
- Device Groups: **DATA PROTECTION > Device Groups**
- File Storage Groups: **SYSTEM > eNAS > File Dashboard > File Storage Groups**

For more information about thin pools and thin provisioning concepts, refer to the Product Guide.

## Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Select the volume type by selecting a tab.
4. Do one of the following:
  - To start space reclamation:
    -  Select one or more volumes, click , and select **Start Allocate/Free/Reclaim** to open the **Start Allocate/Free/Reclaim** dialog box.
    - To start reclaiming unwritten tracks and those written completely with zeros, select **Reclaim Volumes**. This option is only available on storage systems running Enginuity 5876 or higher, or storage systems running HYPERMAX OS 5977 or higher.
    - Optional: To reclaim tracks that are unwritten or zero-based, even if they are marked as persistent, select **Reclaim persistent capacity**.
    - 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.
  - To stop space reclamation:
    -  Select one or more volumes, click , and select **Stop Allocate/Free/Reclaim** to open the **Stop Allocate/Free/Reclaim** dialog box.
    - Select **Stop Reclaim**. In addition, on storage systems running Enginuity 5876, you can optionally specify to free tracks that are unwritten or zero-based, even if they are marked persistent.
    - 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.
    - 
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Advanced Options dialog

Refer to the parent help topic for information on the **Advanced Options** dialog.



# CHAPTER 7

## Performance Management

• Understanding Performance Management.....	518
• Performance Dashboards.....	518
• Charts View.....	527
• Analyze view.....	533
• Heatmap.....	575
• Reports.....	577
• Viewing Real Time traces.....	586
• Plan View.....	588
• Viewing system registrations.....	591
• Viewing Performance databases.....	594
• Viewing Performance thresholds and alerts.....	600
• About exporting and importing performance settings.....	604
• Metrics.....	606
• Viewing and managing metrics.....	718
• Setting the time range for viewing data.....	831
• Creating a template dashboard from an Analyze view.....	831
• Filtering performance data.....	832

# Understanding Performance Management

Performance Management covers the following areas:

- Dashboards - Display pre-defined, user-defined custom dashboards, and templates.
- Charts - Create custom charts across multiple categories/metrics/time intervals.
- Analyze - Provide in-depth drill-down on storage system data for various collection ranges.
- Heatmap - Display hardware instances represented as colored squares, with the color indicating utilization levels.
- Reports - Create, manage and run performance reports.
- Real Time Traces - Create, manage and run performance real time traces (see [Viewing Real Time traces](#) on page 586).
- Databases - Manage Performance database tasks, for example, backup, restore, delete, as well as individual performance database information (see [Viewing Performance databases](#) on page 594).
- Plan - Provide performance projection capacity dashboards displaying predicted future data, based on linear projection.

## Performance Dashboards

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

The Performance component of Unisphere collects performance data only after you [register your storage systems](#).

A dashboard is a collection of charts that you define (custom), or that Unisphere provides. You can view a set of dashboards for a specific storage system (see [Viewing dashboards](#) on page 519), or you can view a set of overview dashboards covering all storage systems (see [Using the All Arrays overview dashboard](#) on page 521).

Custom dashboards must have a unique name up to 32 characters. One or more can be marked as a favorite to allow quick and easy access. Custom dashboards can be promoted to being a favorite by hovering over the dashboard in the table and clicking the Add to Favorites button. The chosen dashboard will then appear in the Favorites selection for easy access.

Only applicable dashboards (those defined in the dashboard catalog - refer to [Managing dashboard catalog](#) on page 594) are displayed for individual storage systems.

The overview dashboards for all storage systems include the following tabs:

- All Arrays Overview
- All Arrays Thresholds

When viewing a specific storage system, the following dashboard folders are provided:

### Favorites

Custom dashboards that have been marked as favorites.

### Default

Pre-configured dashboards that 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.

Depending on the configuration of the individual storage system, all default dashboards might not be available. In addition, you can configure which dashboards are listed in the **Default** folder.

### Custom

Dashboards that were created by saving a chart or a group of charts as a custom dashboard.

### Templates

Dashboards that were created by saving a chart or a group of charts as a template.

The specific information displayed in a default dashboard view depends on the selected storage system and **Category**. Unless the selected category is **Array**, the information is filtered for the selected category (for example, **Database** or **Storage Group**).

Depending on the dashboard displayed, you can perform some of the following tasks:

- Schedule a report to run.
- Navigate directly to the Charts or Analyze view.
- Switch between the use of average and maximum values for the current dashboard.
- Change the number of columns of charts in the current dashboard.
- Change the time frame of the current dashboard.
- Run a PDF report.
- Navigate to the Storage Groups view from the Storage Group dashboard.

To manage dashboards, you can perform the following tasks:

- [Creating a template dashboard from an Analyze view](#) on page 535
- [Creating a dashboard with charts](#) on page 522
- [Copying a dashboard](#) on page 523
- [Saving a dashboard as a template](#) on page 524
- [Viewing dashboards](#) on page 519
- [Using the All Arrays overview dashboard](#) on page 521
- [Deleting dashboards](#) on page 523
- [Saving dashboards and charts](#) on page 526

## Viewing dashboards

### Before you begin

Read [Performance Dashboards](#) on page 518.

To view dashboards:

## Procedure

1. Do one of the following:
  - View a dashboard:
    - a. Select a storage system.
    - b. Select **Performance > Dashboards**.  
To view one of the pre-configured default dashboards, choose one from the **Default** section of the **Dashboards** panel on the left.

---

### Note

You cannot delete default dashboards.






---

To view Custom dashboards, expand the **Custom** dashboards folder.

To view Template dashboards, expand the **Templates** dashboards folder.

- View the overview dashboards for all storage systems:
  - a. Select **Performance > Dashboard**  
By default, the All Arrays Overview dashboard displays.

Depending on the selected dashboard type, some of the following controls are available:

-  **Save As** — [Saving dashboard changes](#) on page 524
-  **Delete Custom Dashboard** or **Delete Template Dashboard**
-  **Schedule PDF Report** — [Scheduling a report from the dashboard](#) on page 525
-  **Run PDF Report now**
  - **Run PDF Report now**
  - **Run PDF Report now - with data**
-  **Navigate to**
  - **Navigate to Charts**
  - **Navigate to Analyze**
  - **Navigate to Details** — [Navigating to the Details view](#) on page 526

## Using default dashboards

### Before you begin

Read [Performance Dashboards](#) on page 518.

To configure what appears in the dashboard, see [Managing dashboard catalog](#) on page 594.

The default dashboards require Administrator, Monitor, or Performance Monitor permissions.

The full list of default dashboards is shown below. A subset of these are presented depending on array type and configuration.



- Array

**Array Efficiency** tab is not applicable for storage systems running Enginuity OS 5876 or non-flash storage systems running HYPERMAX OS 5977.

The Array Performance Thresholds dashboard gives an visual overview to indicate if system component utilization thresholds have been broken for hardware components. Individual bar charts can be double clicked to view more information (including threshold information) (see [Utilization Threshold charts](#) on page 527).

- BE Directors (DA)
- External Directors
- Database
- Disk Groups
- Disk Technology
- FE Directors
- Hosts
- Port Groups
- RDF Directors
- Storage Containers
- Storage Groups
- Disk Group Tier
- Virtual Pool Tier
- Thin Pools

The FE Directors, BE directors (DA), RDF Directors, Disk Groups, Disk Technology, and External Directors have an **All** tab where you can compare all instances of that type that exist on the sotrage system.

### Procedure

1. To save a default dashboard to a PDF file, complete the following steps:


- a. Select a storage system.
- b. Select **Performance > Dashboards**
- c. Expand the **Default** section.
- d. Select a dashboard.

- e.

To run a PDF report, click .

The PDF report is downloaded in the browser.

- f.

Alternatively, to schedule a PDF report, click .

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.`

## Using the All Arrays overview dashboard

### Before you begin

The default dashboards require Administrator, Monitor, or Performance Monitor permissions.

By default, Performance Dashboards across all available arrays are displayed. This allow you to perform top level comparisons across all arrays.

To view the All Arrays overview dashboard, complete the following steps:

#### Procedure

1. Select **Performance > Dashboard**.

A dashboard displaying aggregated data for all arrays displays. The following tabs are displayed:

##### All Arrays Overview

A dashboard displaying the following charts for all arrays:




- Host IOs/sec
- Latency
- Throughput
- Capacity

##### All Arrays Thresholds

A dashboard displaying threshold information for all arrays.

After clicking on an individual array ID, you will be drilled down in context to that's array's full suite of dashboards.


The following controls are available:

-  **Schedule PDF Report** — [Scheduling a report from the dashboard](#) on page 525
-  **Run PDF Report Now**
-  Navigate to
  - **Navigate to Charts**
  - **Navigate to Analysis**

## Creating a dashboard with charts

To create a dashboard with charts:

#### Procedure

1. In the **Charts** view, create one or more charts
2. Click  .
3. In the **Save as Dashboard** dialog box, specify a dashboard name.
4. Click **OK** to save the dashboard.

---

#### Note

If a dashboard with the same name already exists, clicking OK will overwrite that dashboard.

---

The dashboard is saved to the **Custom** dashboard folder.

## Editing a template dashboard

To edit a template dashboard:

### Procedure



1. Select a storage system.
2. Select **Performance > Dashboards**.
3. In the **Templates** list, select the user template dashboard you want to edit.
4. Do one or more of the following:
  - Hover over the template dashboard name in the Templates list and click the star icon to add the template dashboard to the Favorites list.
  - Delete the template dashboard.

## Copying a dashboard

You can copy an existing custom or template dashboard and change the name and some characteristics to make a new dashboard. You cannot copy a default dashboard.

To copy a dashboard:

### Procedure

1. Select a storage system.
2. Click **Performance > Dashboards**
3.  Select the dashboard to copy and click  to open the **Copy Dashboard** dialog box.
4. Type a **Dashboard Name** for the copied dashboard.
5. Click **OK**.

## Editing dashboards

To edit a dashboard:

### Procedure

1. Select a storage system.
2. Select **Performance > Dashboards** to view the list of dashboards.
3. Select the dashboard you want to edit.
4. Do one or more of the following:
  - Hover over the template dashboard name in the Templates list and click the star icon to add the template dashboard to the Favorites list.
  - Delete the template dashboard.


## Deleting dashboards

### Before you begin

You can delete only dashboards in the Custom or Templates folders. Predefined dashboards cannot be removed.

To delete a dashboard:

**Procedure**

1. Select a storage system.
2. Select **Performance > Dashboards**.
3. In the Custom or Templates folder, hover over the dashboard you want to delete and click  .
4. Click **OK** on the delete confirmation message.


**Running a report from the dashboard**

You can run a PDF report from a performance dashboard.

**Procedure**

1. Select a storage system.
2. Click **Performance > Dashboards**.
3. Open an existing dashboard or create a new one.
- 4.




Click  to run the report now.

**Saving a dashboard as a template**

Selected default dashboards can be adjusted and saved as a template for your own dashboard as long as they are for the same storage system, category, and object type as the original.

To save a dashboard as a template:

**Procedure**

1. Select a storage system.
2. Select **Performance > Dashboards**.
3. Select one of the default dashboards.
4. Click  > **Save as Template** to open the **Save as Template** dialog box.
5. Enter a **Template Name**.  
The template is saved to the **Templates** folder.
6. Click **OK**.

**Saving dashboard changes**

The **Save as** option allows you to do the following:

- Save charts created in the chart section as either a Custom Dashboard or a Template.
- Save an existing Default or Custom Dashboard as a Template.
- Save a Template as a Custom Dashboard.

**Note**

You cannot save a dashboard containing multiple tabs or item lists.

To save dashboard changes:

### Procedure

1. Click **Performance** > **Dashboards** to view the list of dashboards.
2. Open the dashboard chart that you want to change.
3. Update the chart, depending on the options available.

For more information, see [Customizing the tabbed Charts view](#) on page 530 and [Customizing a chart](#) on page 528.


4. Click  > **Save as Custom Dashboard**.

## Scheduling a report from the dashboard

The performance dashboard reports are saved in PDF format. When you schedule a report, it will run on the time/days selected.

To schedule a report from the dashboard:

### Procedure

1. Select a storage system.
2. Select **Performance** > **Dashboards**.
3. Select a default or custom dashboard.
4. Click  to open the **Create Report** dialog box.
5. In the **General** section, enter a report **Name**.
6. (Optional) Enter the report description.
7. In the **Schedule** section, select **Schedule**.
8. Set **First Run Time**.

The **First Run Time** is continually adjusted relative to the current time and the report execution interval. For example, if **First Run Time** is set to 3am, but the current time is 4am and the **Time Range** is set to 'Last 4 hours', then the value of **First Run Time** is changed to 7am.

9. Select the days on which the report will run.
10. From the **Retention (Days)** menu, select the number of days for which you want the generated reports to be retained.
11. (Optional) In the **Email** section, do one or more of the following:
  - Select the **Email report to** check box to send an email to a preconfigured email address.
  - Select the **Send report to** check box and type a valid email address in the field.
12. Click **OK**.


A confirmation message displays. Scheduled reports are listed under **Performance** > **Reports**. Changes can be made to scheduled reports from there.

Scheduled reports save to this default location:

```
install_dir\SMAS\report\SPA\querydata\
```


## Navigating to the Details view

From many of the Dashboards you can navigate to the component's details view.

If the  icon is dimmed, there is no corresponding detail view.

To navigate to the details view:

### Procedure

1. Select a storage system and an existing dashboard.
2. Click .

The component's details display. For example, the details view for the storage group, front-end director, or disk group.

You cannot navigate from any of the dashboards for all components (all arrays, all disk groups, and so on.).

## 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 (JPEG or PNG file format). When exporting to PDF, you can export the chart or dashboard only, or the dashboard or 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 or chart.

### Saving a dashboard or chart to an image file

#### Procedure

1. Open the dashboard or chart that you want to save.
- 2.


In the individual chart, click  > **Save As**.

3. Select the image format you want to use. The following choices are available:
  - JPEG
  - PNG
  - CSV

The image is downloaded in the browser.

### Saving a dashboard to PDF

#### Procedure

1. Open the dashboard you want to save.
2. Click .
3. The PDF copy of the dashboard is downloaded in the browser.

## Opening a dashboard image in a new browser tab

### Note

This functionality is available only when using Firefox or Chrome web browsers.

### Procedure

1. Using Firefox or Chrome, open the dashboard or chart that you want to save.
- 2.

In the individual chart, click  > **Save As > To Clipboard**.

3. A JPG image of the current dashboard opens 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.

## Utilization Threshold charts

The Array Performance Thresholds dashboard gives an visual overview to indicate if system component utilization thresholds have been broken for hardware components.

Individual bar charts can be double clicked to view more information (including threshold information).

## Charts View

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

The All Array Charts view allows you to easily draw charts for any performance registered array without having to drill down to the performance section for an individual array.

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:

### KPI Metrics

Lists the key performance indicators that correspond to the column headings in the current table.

### All Metrics

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 maximize chart viewing space the New Chart panel can be collapsed by clicking the Close Sidebar button. To restore the New Chart Panel click the Add New Chart button.

To create a chart:

### Procedure

1. Select the storage system.

2. Select **Performance > Charts > New Chart**.
3. (Optional) In the Chart Details section, type a **Chart Title**.
4. Select an appropriate time range for the chart or specify a custom time range.
5. In the Objects section, select the systems and/or component types for which you want to create a chart.
6. (Optional) If required, move the slider from **Avg** to **Max** to use maximum metric values rather than average values, which are used by default.
7. (Optional) If required, click  to filter the list of selected components by name.
8. You can sort the list alphabetically by clicking on the object name at the top of the list.
9. Select the component instances. Multi-select is available.
10. (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.
11. Select the one or more metrics from the **Metrics** list or use the slider to display the **All Metrics** list and select metrics from that list. You can filter the list by clicking on  and typing a metric name in the search box.  
You can select multiple metrics, however, the more you select, the less readable the chart will be.
12. Click **Create**.  
Depending on the selected options, you might be given a choice between creating different types of charts. The options available are:
  - **One Chart**
  - **One Chart per Instance**
  - **One Chart per Metric**
 Select an appropriate option and click **Create**.
13. (Optional) Customize the chart. For more information about how to do this, refer to [Customizing a chart](#) on page 528.
14. (Optional) Customize the Chart view in which the chart is displayed. For more information about how to do this, refer to [Customizing the tabbed Charts view](#) on page 530.

## Customizing a chart

After you create a chart there several ways to customize that chart.

For more information about editing the information displayed on a chart, see [Editing charts](#) on page 532.

For more information about customizing the tabbed Charts view, see [Customizing the tabbed Charts view](#) on page 530.

### Procedure

1. To customize a chart, hover over the chart, and do one or more of the following, depending on the options available:
  - Change the chart type:



a.

Click **Chart Type**. The following options are available:

- **Timeline**
- **Snapshot**
- **Threshold Table**

Save the chart as an image:

a.

Click **> Save as**.

- b. Select JPG, PNG or CSV.
- c. Specify a file name for the image file.

- Change the style of the chart:

a.

Click **> Chart Style**.

- b. Select a chart style.

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:

- **Line**
- **Bar**
- **Area**
- **Table**

- Change the data format of the chart:

a.

Click **Data Format**. Select one or both of the following:

- **Average** — Shows the average metric value for each metric in the table rows, for the session (hour). This is the default.
- **Max** — Shows the maximum metric value—the high water mark—for each metric in the table rows, for the session (hour).

Selecting both formats on a chart displaying multiple metrics can make the chart less readable.

- Save the chart as a JPEG file:

a.

Click **> Save as > JPEG**

- b. Specify a file name for the JPEG file.

- Save the chart as a PNG file:

a.

Click **> Save as > PNG**

- b. Specify a file name for the PNG file.

- Save the chart as a CSV file:

a.

Click **> Save as > CSV**

- b. Specify a file name for the CSV file.
- Display a threshold on the chart (if thresholds have been set for a metric):
  - a.









Click  > **Display Thresholds.**

---

#### Note

Threshold values are not displayed in table charts.

---

- Display service level information on the chart:
  - a.
- Click  > **Display Service Level.**
- Hide the chart legend:
  - a.
- Click  > **Hide Legend.**
- Single Scale
  - a.
- Click  > **Single Scale.**
- Show information about the chart:
  - a.
- Click  .
- b. To close the information pane, click .
- Maximize the chart:
  - a.
- Click  .
- Minimize the chart:
  - a.
- Click  .
- Restore a maximized chart to the default chart size:
  - a.
- Click  .
- Remove the chart:
  - a.
- Click  .
- Change the zoom level of a chart:
  - a. Select an area on a chart to zoom into that area.
  - b. To reset the zoom level to default, click **Reset zoom.**

## Customizing the tabbed Charts view








There several ways to customize the way that charts are displayed on the tabbed Charts view.

For more information about customizing the way that individual charts are displayed, see [Customizing a chart](#) on page 528.


For more information about editing the information displayed on a chart, see [Editing charts](#) on page 532.


Tabs can be closed by selecting the "X" button on the tab. There must be at least one tab available.


### Procedure


1. When in the tabbed Charts view, and without any individual charts selected, do one, or more, of the following:
  - Add a chart:
    - a. If the **New Chart** menu is not already open, in the Charts view click **Add Charts**.
    - b. Create a new chart. For more information, refer to [Charts View](#) on page 527.
  - Create a new tab for charts:
    - a. Click  .
  - Save the Charts view as a dashboard:
    - a. Click  .
    - b. In the **Save as Dashboard** dialog box, type a **Dashboard Name**.
    - c. Click **OK**.
  - Save the Charts view as a template (Note that all Charts in the tab must all be for the same category, Object and time period.):
    - a. Click  > **Save As Template**
    - b. In the **Save as Template** dialog box, type a **Template Name**.
    - c. Click **OK**.
  - Save the Charts view as a PDF:
    - a. Click  > **Save As PDF - Charts** .
  - Save the Charts view as a PDF with data included in an additional table:
    - a. Click  > **Save As PDF - Charts & Data** .
  - Remove all charts from the Charts view:
    - a. Click  to select all charts.
    - b. Click  .
  - Adjust the number of columns displayed in the Chart view:
    - a. From the **Chart Column(s)** drop-down menu, select a value.
2. When in the tabbed Charts view, and with an individual chart selected (select the chart by selecting the checkbox in the top left corner of the chart), do one, or more, of the following:


- Copy a chart to another tab:
  - a.
 


Click  > **Copy Charts to tab**
  - b. In the **Copy charts to existing tab** dialog box, select a tab.
  - c. Click **OK**.
- Move a chart to another tab:
  - a.
 

Click  > **Move Charts to tab**
  - b. In the **Move charts to existing tab** dialog box, select a tab.
  - c. Click **OK**.
- Save the Charts view as a template (Note that all Charts in the tab must all be for the same category, Object and time period.):
  - a.
 

Click  > **Save As Template**
  - b. In the **Save as Template** dialog box, type a **Template Name**.
  - c. Click **OK**.
- Save the Charts view as a PDF:
  - a.
 

Click  > **Save As PDF - Charts .**
- Save the Charts view as a PDF with data included in an additional table:
  - a.
 

Click  > **Save As PDF - Charts & Data .**
- Remove the chart from the Charts view:
  - a.
 

Click  .
- Adjust the number of columns displayed in the Chart view:
  - a. From the **Chart Column(s)** drop-down menu, select a value.

## Editing charts


You can edit the information displayed on a chart.

For more information about customizing the way that individual charts are displayed, see [Customizing a chart](#) on page 528.

For more information about customizing the tabbed Charts view, see [Customizing the tabbed Charts view](#) on page 530.

### Procedure

1. Select the chart using the checkbox in the top left corner.
2.
 

Click  to open the **Edit Chart** pane.
3. Do one of the following:

- Add a metric:
  - a. Select objects and apply filters, as appropriate to display the list of metrics.
  - b. Select the metric in the **KPI Metrics** or **All Metrics** list.
  - c. Click **Save**.
- Remove a metric:
  - a. Clear the checkbox alongside the metric you want to remove from the chart.
  - b. Click **Save**.

## Copying a chart

To copy a chart:

### Procedure

1. Select a storage system.
2. Click **Performance > Charts**.
3. Select the chart by clicking the checkbox in the top left corner of the chart and

then click  to create the copy.

## Analyze view

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

Only applicable dashboards (those defined in the dashboard catalog - refer to [Managing dashboard catalog](#) on page 594) are displayed for individual storage systems.

The All Array Analyze view allows you to easily compare and drill into analyze data for any performance registered array without having to drill down to the performance section for an individual array.

The Analyze view provides three tabbed views into your storage data. You can use the Real Time and Diagnostic tabs to debug and find causes for performance problems. You can use the Historical tab to view historical data, showing long-term trends, to plan for future needs.

### Real Time tab

The default range is the previous 1 hour. Real Time view collects data between 2 and 5 seconds for a limited group of metrics. The data is available for the previous hour. Saved Real Time traces can also be viewed here.

### Diagnostic tab

The default range is the previous 4 hours. Diagnostic view collects data every 5 minutes (by default, but this can be changed to 10 or 15 minutes) for root cause analysis. The data is available for the previous 15 days (by default, but this can be reduced). A maximum of a 24 hour time period can be viewed at a time.

### Historical tab

The default range is the previous 24 hours. Historical view collects data in hourly or daily intervals for trending and planning. The data is available for a maximum of the previous 3 years (18 months by default).

You can use Real Time and Diagnostic information to find the cause of performance problems. Historical information shows long-term trends you can use to plan for future requirements.

### Procedure

1. Select a storage system.
2. Click **Performance > Analyze**
3. Select a tabbed view. The following tabs are available:
  - Real Time
  - Diagnostic
  - Historical



The Diagnostic tab is displayed by default.

4. To see more information about a storage system, click the ID of the storage system.

More detailed data, displayed according to category, is presented.

5. To navigate between the different categories of information for that storage system, select 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.

6. (Optional) To configure the columns presented click  and select the columns you want to view. For some analyze views, additional columns can be viewed.
7. To filter a column, click  and type letters or numbers in the search box, depending on the data type in the column being searched. For columns with numerical values you can filter for exact values or for ones larger than or equal or smaller than or equal to the value e.g. 100, >100, >=100, <100, <=100. You can also filter for a numerical range e.g. 100-150.
8. (Optional) To return to the top-level, table view of storage systems, click the storage system ID in the breadcrumb trail.
9. (Optional) To take a trace for the last one hour's worth of real time data. click the **Capture Trace** button (available from the **Real Time** tab).

## Creating a dashboard from an Analyze view


### Before you begin

You require access to a visible, registered storage system.


You can create a dashboard from the information currently displayed in an Analyze view.

To create a dashboard from an Analyze view:

### Procedure

1. Select component(s) from the table in an Analyze view.
2. Click  > **Navigate to Charts**

The selected components are displayed in a **New Chart** view.

3. Select the metrics to be displayed.
4. Click **Create**.
5. Click .


The **Save as Dashboard** dialog box displays.

6. Specify a dashboard name.
7. Click **OK** to save the dashboard.

## Creating a template dashboard from an Analyze view

To create a dashboard template from an Analyze view:

### Procedure

1. Select component(s) from the table in an Analyze view.
2. Click  > **Navigate to Charts**

The selected components are displayed in a **New Chart** view.

3. Select the metrics to be displayed.
4. Click **Create**.
- 5.

Click  > **Save as Template**.

The **Save as Template** dialog box displays.

6. Specify a template name.
7. Click **OK** to save the template.

## 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** view, there are the following default data collection time ranges:

### Real Time

The last hour.

### Diagnostic

The last 4 hours.

### Historical

The last 24 hours.

For the **Real time** tab, you can load an existing real-time trace file by selecting **Realtime** > **Trace** from the time-selection drop-down menu.

For the **Diagnostic** and **Historical** tabs, you can customize the time range by specifying a custom time range from the time-selection drop-down menu.

When specifying a custom Diagnostic time range, you can define a maximum range of 24 hours, specified at five-minute intervals.

When specifying a custom Historical time range, you can define a minimum range of 24 hours, specified at one-hour intervals.

## Procedure

### 1. Complete one of the following:

- Changing the Real Time time range:

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 available only if a trace file is available for the selected storage system.

- Changing the Diagnostic time range:

a. From the time selection menu, which by default displays `Last 4 Hours`, select **Diagnostic**, and 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:

a. In the **Custom Time Selection** dialog box, specify a start time and end time.

b. Click **OK**.

- Changing the Historical time range:

a. From the time selection menu, which by default displays `Last 24 Hours`, select one of the following options:

- Select **Last 24 Hours** to view the last 24 hours of data collected.
- Select **Last Week** to view the last week of data collected.
- Select **Last Month** to view the last month of data collected.
- Select **Last 6 Months** to view the last 6 months of data collected.
- Select **Last 12 Months** to view the last 12 months of data collected.
- Select **Custom** to make a custom time selection. To do this, complete the following steps:

a. In the **Custom Time Selection** dialog box, specify a start time and end time.

b. Click **OK**.

## Symmetrix systems view (Real Time)

The Real Time view provides high-level data about all storage systems. The specific information displayed depends on the models of the selected storage systems and consists of some or all of the following:

### 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)**

The Real Time view provides detailed data about front-end directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**Name**

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)**

The Real Time view provides detailed data about back-end directors (DA). The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**Name**

The ID of the back-end 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)

The Real Time view provides detailed data about external directors. The specific information displayed depends on the model of the selected storage system and consists of Some or all of the following:

**Name**

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)

The Real Time view provides detailed data about RDF directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**Name**

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.

## Array systems view (Diagnostic)

The Diagnostic view provides high-level data about all storage systems. The specific information displayed depends on the models of the selected storage systems and consists of some or all of the following:

### **Name**

The array 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.

## Alerts view (Diagnostic)

The Diagnostic view provides detailed data about alerts. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### **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, and so on.

**Instance**

The specific object on which the alert is generated. 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 (Diagnostic)

The Diagnostic view provides detailed data about RDF directors. The specific information displayed depends on the model of the selected storage system, and consists of some or all of the following:

**Name**

The ID of the director.

**% Busy**

The percent of time the director 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 (Diagnostic)**

The Diagnostic view provides detailed data about back-end directors (DA). The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**Name**

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 (Diagnostic)**

The Diagnostic view provides detailed data about external directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**Name**

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 (Diagnostic)

The Diagnostic view provides detailed data about RDF directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**Name**

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 (Diagnostic)

The Diagnostic 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:

**Name**

The ID of the IM director.

**% Busy**

The percent of time that a director is busy.

**% Idle**

The percent of time that a director is idle.

## EDS Directors view (Diagnostic)

The Diagnostic view provides detailed data about EDS directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### **Name**

The ID of each EDS director.

### **% Busy**

The percent of the time that a director is busy.

### **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 (Diagnostic)

The Diagnostic view provides detailed data about cache partitions. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### **Name**

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 (Diagnostic)

The Diagnostic view provides detailed data about boards. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### **ID**

The ID of the board.

### **Utilization**

The percent of time the board is busy.

**Outbound MBs/sec**

-

**Inbound MBs/sec**

-

**Max Speed MBs/sec**

-

## Disk Technologies view (Diagnostic)

The Diagnostic view provides detailed data about disk technologies. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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.

**Read RT (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.

**Write RT (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 (Diagnostic)

The Diagnostic view provides detailed data about events. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:



**Sequence Number**

-

**Timestamp**

-

**Function Class**

-

**Action Code**

-

**Activity ID**

-

**Message**

-

**Application ID**

-

**Host ID**

-

**Vendor ID**

-

**User name**

-

**Offset**

-

## Storage Groups view (Diagnostic)

The Diagnostic view provides detailed data about storage groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 RT (ms)**

The time it takes the group to perform the Reads in milliseconds.

**Write RT (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.

**Allocated Capacity (GB)**

The total allocated array capacity.

**Total Capacity (GB)**

The capacity of the storage group in GBs.

**Members**

The number of volumes that comprise this storage group.

## Device Groups view (Diagnostic)

The Diagnostic view provides detailed data about device groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

The name assigned to the device group.

**Host IOs/sec**

The number of host operations performed each second by the group.

**Read RT (ms)**

The average time it took the storage system to serve one read IO for this group.

**Write RT (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.

**Allocated Capacity**

The total allocated array capacity.

**Total Capacity (GB)**

The capacity of the device group in GBs.

**Members**

The number of volumes that comprise this device group.

**Databases view (Diagnostic)**

The Diagnostic view provides detailed data about databases. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 (Diagnostic)**

The Diagnostic view provides detailed data about virtual pools. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 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.

**Total Capacity**

The total capacity of the thin pools in GBs.

## Disk Groups view (Diagnostic)

The Diagnostic view provides detailed data about disk groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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.

**Read RT (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.

**Write RT (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 (Diagnostic)**

The Diagnostic view provides detailed data about external disk groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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.

**Read RT (ms)**

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.

**Write RT (ms)**

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 (Diagnostic)**

The Diagnostic view provides detailed data about SRPs. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

The ID of the SRP.

**Host IOs/sec**

The number of IOs per second transferred from the host.

**Host MBs/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.

**Response Time (ms)**

The average time it takes the SRP to service IO.

## **RDFA Groups view (Diagnostic)**

The Diagnostic view provides detailed data about RDF/A groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

Displays the RDF/A group number.

**RDFA Group Label**

Displays the RDF/A group name.

**RDF Response Time (ms)****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 (Diagnostic)**

The Diagnostic view provides detailed data about RDF/S groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

The RDF/S group number.

**RDF/S Group Label**

The RDF/S group label.

**Response Time (ms)**

The average time it took the storage system to serve one IO for the RDF/S group.

**Host Writes/sec**

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.

**% Writes**

The percent of IO operations that were writes for the RDF/S group.

**Snap Pools view (Diagnostic)**

The Diagnostic view provides detailed data about snap pools. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 RT (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 (Diagnostic)**

The Diagnostic view provides detailed data about DSE pools. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 RT (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 VP Policies view (Diagnostic)**

The Diagnostic view provides detailed data about FAST policies. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 (Diagnostic)**

The Diagnostic view provides detailed data about disk group tiers. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:



**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.

**Read RT (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.

**Write RT (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 (Diagnostic)

The Diagnostic view provides detailed data about virtual pool tiers. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 RT (ms)**

The average time it takes to perform IO requests from the disk directors to cache.

**Allocated Pool capacity (GB)**

The allocated pool capacity.

**Total Capacity (GB)**

The total pool capacity.

## Storage Groups view (Diagnostic)

The Diagnostic view provides detailed data about storage groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 RT (ms)**

The time it takes the group to perform the Reads in milliseconds.

**Write RT (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.

**Allocated Capacity (GB)**

The total allocated array capacity.

**Total Capacity (GB)**

The capacity of the storage group in GBs.

**Members**

The number of volumes that comprise this storage group.

## Hosts view (Diagnostic)

The Diagnostic view provides detailed data about hosts. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### 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.

### Response Time (ms)

The overall average response time for reads and writes.

### 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.

### Host Reads/sec

Cumulative number of host MBs read per second.

### Host Writes/sec

Cumulative number of host MBs written per second.

### Syscall Count/sec

The total number of calls seen by this host.

## Initiators view (Diagnostic)

The Diagnostic view provides detailed data about initiators. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### 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 (Diagnostic)

The Diagnostic view provides detailed data about masking views. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### 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 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.

### Capacity (GB)

The total capacity of the masking group, in GBs.

## Port Groups view (Diagnostic)

The Diagnostic view provides detailed data about port groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### 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 by SG view (Diagnostic)

The Diagnostic 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 and consists of some or all of the following:

### ID

The name of the storage group.

### Host IOs/sec

Number of IOs the FE port is completing for the storage group.

**Read RT (ms)**

The average time it takes to serve one read IO for this storage group.

**Write RT (ms)**

The average time it takes to serve one write IO for this storage group.

**Host IO Limit IOs/sec**

The number of IOs the FE port is processing on behalf of the specific storage group.

**Host IO Limit MBs/sec**

The number of MBs the FE port is processing on behalf of the specific storage group.

**Host IO Limit Exceeded secs**

The number of IOs the FE port is performing on behalf of the specific storage group.

**Host IO Limit by FE view (Diagnostic)**

The Diagnostic 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 and consists of some or all of the following:

**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 Exceeded Secs**

The number of IOs the FE port is performing on behalf of the specific SG.

**Symmetrix systems view (Historical)**

The Historical view provides high-level data about all storage systems. The specific information displayed depends on the models of the selected storage systems and consists of some or all of the following:

**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.

## Alerts view (Historical)

The Historical view provides detailed data about alerts. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 (Historical)

The Historical view provides detailed data about front-end directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 (Historical)

The Historical view provides detailed data about back-end directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 (Historical)

The Historical view provides detailed data about external directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 (Historical)

The Historical view provides detailed data about RDF directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 (Historical)**

The Historical view provides detailed data about IM directors. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

The ID of the IM director.

**% Busy**

The amount of time that the director is busy.

**% idle**

The amount of time that the director is idle.

**EDS Directors view (Historical)**

The Historical view provides information about EDS directors:

**ID**

The ID of each EDS director.

**% Busy**

The amount of time that the director is busy.

**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 (Historical)**

The Historical view provides detailed data about cache partitions. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

#### **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 (Historical)**

The Historical view provides detailed data about boards. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

#### **ID**

The ID of the board.

#### **Utilization**

A measure of how busy the board is calculated as the sum of the input and output transfer rates divided by the maximum transfer rate.

#### **Outbound MBs/sec**

-

#### **Inbound MBs/sec**

-

#### **Max Speed MBs/sec**

-

## **Disk Technologies view (Historical)**

The Historical view provides detailed data about disk technologies. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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.

**Read RT (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.

**Write RT (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 (Historical)

The Historical view provides detailed data about events. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**Sequence Number**

-

**Timestamp**

-

**Function Class**

-

**Action Code**

-

**Activity ID**

-

**Message**

-

**Application ID**

-

**Host ID**

-

**Vendor ID**

-

**User name**

-

**Offset**

-

## Storage Groups view (Historical)

The Historical view provides detailed data about storage groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

The name assigned to this storage group.

**Service Level**

The service level associated with the 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 RT (ms)**

The time it takes the group to perform the Reads in milliseconds.

**Write RT (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 (Historical)**

The Historical view provides detailed data about device groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

The name assigned to the device group.

**Host IOs/sec**

The number of host operations performed each second by the group.

**Read RT (ms)**

The average time it took the storage system to serve one read IO for this group.

**Write RT (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.

**Allocated Capacity (GB)**

The allocated capacity of the device group in GBs.

**Total Capacity (GB)**

The capacity of the device group in GBs.

**Members**

The number of volumes that comprise this device group.

**Databases view (Historical)**

The Historical view provides detailed data about databases. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 (Historical)

The Historical view provides detailed data about virtual pools. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 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.

**Total Capacity (GB)**

The total capacity of the thin pool in GBs.

**Disk Groups view (Historical)**

The Historical view provides detailed data about disk groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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.

**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.

**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 (Historical)**

The Historical view provides detailed data about external disk groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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.

**Read RT (ms)**

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.

**Write RT (msec)**

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 (Historical)

The Historical view provides detailed data about SRPs. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

The ID 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.

**Response Time (ms)**

The average time it takes the SRP to service IO.

## RDFA Groups view (Historical)

The Historical view provides detailed data about RDF/A groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:



**ID**

Displays the RDF/A group number.

**RDF/A Group Label**

Displays the RDF/A group name.

**RDF Response Time (ms)**

The total writes 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 (Historical)

The Historical view provides detailed data about RDF/S groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

The RDF/S group number.

**RDF/S Group Label**

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.

**% Writes**

The percent of IO operations that were writes for the RDF/S group.

## Snap Pools view (Historical)

The Historical view provides detailed data about snap pools. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### 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 RT (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 (Historical)

The Historical view provides detailed data about DSE pools. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

### 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 RT (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 VP Policies view (Historical)**

The Historical view provides detailed data about FAST policies. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 (Historical)**

The Historical view provides detailed data about disk group tiers. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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.

**Read RT (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.

**Write RT (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 (Historical)

The Historical view provides detailed data about virtual pool tiers. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 Response Time (ms)**

The average time it takes to perform IO requests from the disk directors to cache.

**Allocated Capacity (GB)**

The allocated pool capacity.

**Total Capacity (GB)**

The total pool capacity.

## Storage Groups view (Historical)

The Historical view provides detailed data about storage groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 view (Historical)

The Historical 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 (Historical)

The Historical view provides detailed data about masking views. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 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.

**Capacity (GB)**

The total capacity of the masking group, in GBs.

## Port Groups view (Historical)

The Historical view provides detailed data about port groups. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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 by SG view (Historical)

The Historical view provides detailed data about host IO limits set for the storage group (SG). The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**ID**

The name of the storage group.

**Host IOs/sec**

Number of IOs the FE port is completing for the storage group.

**Read RT (ms)**

The average time it takes to serve one read IO for this storage group.

**Write RT (ms)**

The average time it takes to serve one write IO for this storage group.

**Host IO Limit IOs/se**

The number of IOs that the FE port is processing on behalf of the SG.

**Host IO Limit MBs/sec**

The number of MBs that the FE port is processing on behalf of the SG.

**Host IO Limit % Time Exceeded****Host IO Limit by FE view (Historical)**

The Historical 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 and consists of some or all of the following:

**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 Exceeded Secs****Initiators view (Historical)**

The Historical view provides detailed data about initiators. The specific information displayed depends on the model of the selected storage system and consists of some or all of the following:

**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.

**Heatmap**







See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

The heatmap displays hardware instances of the selected array represented as colored tiles, with the color indicating utilization levels.

The specific information displayed in a heatmap depends on the model of the selected storage system and whether any filters have been applied. Gray tiles denote components that are not in use.

A graphic indicating the colors used to denote utilization levels is displayed beneath the heatmap.

### Procedure

1. Select the storage system.
2. Select **Performance > Heatmap**.
3. Do one of the following:
  - Hover over a colored tile to instance to view popup utilization information for that hardware instance.
  - Click on a colored tile to open the **Heatmap Metrics Chart(s)** for the hardware instance (see [Viewing Heatmap Metrics Charts](#) on page 576) .
  - If the array has multiple engines, scroll through the heatmaps for them using the chevrons on the left and right.
  -  Click  to export the heatmap to a PDF file that is downloaded by your browser.
  -  Click  to create a report based on the heatmap. For more information, see [Creating performance reports](#) on page 579.
  -  Click to  to **Navigate to Charts**, **Navigate to Analyze**, or **Navigate to Details**, depending on availability. For more information, see [Navigating from Heatmap to Analyze or Charts](#) on page 577.
  - Click the slider to choose between average and maximum metric values when drawing the heatmap.
  - Select a time period from the drop-down list to redraw the heatmap for a different time period.

---

### Note

To produce a future projection heatmap, you need a minimum of ten days of data.

---

- Apply a filter. For more information, see [Filtering heatmaps](#) on page 577.

## Viewing Heatmap Metrics Charts

The heatmap displays hardware instances of the selected array represented as colored tiles, with the color indicating utilization levels.

### Procedure

1. Select the storage system.
2. Select **Performance > Heatmap**.
3. Click on a colored tile to open the **Heatmap Metrics Chart(s)** for the hardware instance.





4. Optional : Click **NAVIGATE TO ANALYZE**.
5. Optional : Click **NAVIGATE TO CHARTS**.
6. Click **OK**.

## Navigating from Heatmap to Analyze or Charts

You can navigate from the heatmap object to a charts or analyze view.


To navigate from Heatmap to Analyze or Charts:

### Procedure

1. Select the storage system.
2. Select **Performance > Heatmap**.
3. Do one of the following:
  - a.  Click .
  - b. From the drop-down menu, select one of the following:
    - **Navigate to Charts**
    - **Navigate to Analyze**
  - a. Click on any object in the heatmap to open its corresponding metric chart.
  - b. Click **Navigate to Analyze** or **Navigate to Charts**.

## Filtering heatmaps

### Procedure

1. Click the filter icon  to open the filter dialog box.
2. From the **Context** menu, select an object type. The instances list updates to show the objects of that type.
3. Select an object from the instances list. You can filter the instances list by typing all or part of the object's name in the search field.
4. Click **OK**.

## Reports

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

To view performance reports:

### Procedure

1. Select the storage array.
2. Select **Performance > Reports**.

Use the Reports list to display and manage reports.

The following properties display:

#### Name

The user-defined report name.

**Array**

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.

**Email Enabled**

Indicates if email notifications are enabled.

**Description**

The user-defined description.

The following controls are available:

- **Create** — [Creating performance reports](#) on page 579
- **Modify** — [Modifying performance reports](#) on page 583
- **Delete** — [Deleting performance reports](#) on page 584
- **Run Now** — [Running performance reports](#) on page 584
- **Schedule** — [Scheduling performance reports](#) on page 584
- **Cancel Schedule** — [Cancelling a scheduled report](#) on page 585
- **Copy** — [Copying performance reports](#) on page 580

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
```

**Note**

If the report contains multiple queries involving volumes and disks the report is saved to this location:

```
/SMAS/report/SPA/runnow
```

Scheduled UPV reports are saved to this location:

```
SMAS\backup\SPA
```

## Report operations

The following report operations are available:

- Create Compliance reports (see [Creating Compliance Reports](#) on page 157).

- Create performance reports (see [Creating performance reports](#) on page 579).
- Modify performance reports (see [Modifying performance reports](#) on page 583).
- Copy performance reports (see [Copying performance reports](#) on page 580).

## Creating performance reports

To create a performance report:

### Procedure

1. Select a storage array.
2. Select **Performance > Reports**.
3. Click **Create**.
4. In the **General** tab, enter a report **Name**.
5. (Optional) Enter a report **Description**.
6. Select the array.
7. Select the **Time Frame**.

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 **Diagnostic** type. This option executes a report on the array and director categories at each Diagnostic interval. In addition, every hour the system runs a cleanup job to remove Last Interval reports that are older than one hour.

8. Select a **Format**

XML, CSV, BTP and UPV formats are supported. If you choose **Last Hour (Realtime)** as the time frame the only format available is CSV. BTP reports are not available for storage systems running HYPERMAX OS 5977 or higher. The BTP format is also not available for **Last Interval** reports.

9. From the **Generated Time Zone** menu, select the time zone to apply to times in the report. (Not applicable to BTP reports).
10. (Optional) For UPV reports, specify whether to include RT traces.
11. 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](#) on page 580. Queries can also be edited (see [Editing queries using the Edit Query wizard](#) on page 582) or deleted.

---

### Note

If the report contains multiple queries involving volumes and disks the report is saved to this location:

```
/SMAS/report/SPA/runnow
```

---

12. (Optional) In the **Schedule** tab, select the **Schedule** check box and configure values for each of the following options:
  - **First Run Time**  
The **First Run Time** is continually adjusted relative to the current time and the report execution interval. For example, if **First Run Time** is set to 3am,

but the current time is 4am and the **Time Range** is set to 'Last 4 hours', then the value of **First Run Time** is changed to 7am.

- **Day(s) to run**
- **Retention**

13. (Optional) In the **Email** tab, specify the report email options.
14. Click **OK**.

## Copying performance reports


### Before you begin

Before copying a performance report, you require an existing report in the **Reports** list view.

To copy a performance report:

### Procedure

1. Select the storage array.
2. Select **Performance > Reports** to open the **Reports** list view.
- 3.

Select a report and click  > **Copy**.

4. Update the report settings, as appropriate. You can edit the following data:
  - Name
  - Description
  - Symmetrix ID
  - Format
  - Generated Time Zone
  - TimeFrame
5. In the **Queries** section, do one of the following:
  - Click **Add** to open the **Create Query** wizard and create a query.
  - Select a query and click **Delete**. Click **OK** on the confirmation message.
  - Select a query and click **Edit** to open the **Edit Query** wizard.  
A report must have at least one query.
6. (Optional) In the **Schedule** tab, configure the schedule settings for the report.
7. (Optional) In the **Email** tab, configure the email settings for the report.
8. Click **OK**.

For more information about reports, see [Creating performance reports](#) on page 579

## 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 page 579.

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:

#### Procedure

1. When creating a report, in the **Create Report** dialog box, click **Add** to open the **Create Query** wizard.
2. In the **Query Name & Description** page, enter a query **Name**.
3. (Optional) Enter a query **Description**.
4. Click **Next**.
5. In the **Category & Instance** page, 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 Instances** list and click > to move them to the **Selected Instances** list. You can use multi-select for multiple instances. If required, you can filter the instances by name.

#### 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**.
8. In the **Metrics** page, select a **Data Format**. The available 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. If required, you can filter the metrics by name.
10. Click **OK**.

## Performance report query operations

- Create query (see [Creating queries using the Create Query wizard](#) on page 580).
- Edit query (see [Editing queries using the Edit Query wizard](#) on page 582).

## Editing queries using the Edit Query wizard

### Before you begin

The **Edit Query** wizard is launched from the **Create Report** dialog box. For more information about creating a report, see [Creating performance reports](#) on page 579.

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 edit a query using the Create Query wizard:

### Procedure

1. When creating a report, in the **Create Report** dialog box, click **Edit** to open the **Edit Query** wizard.
2. In the **Query Name & Description** page, modify the query **Name**.
3. (Optional) Modify the query **Description**.
4. Click **Next**.
5. In the **Category & Instance** page, 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 Instances** list and click > to move them to the **Selected Instances** list. You can use multi-select for multiple instances. If required, you can filter the instances by name.

**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**.

8. In the **Metrics** page, select a **Data Format**. The available 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. If required, you can filter the metrics by name.

10. Click **OK**.

## Modifying performance reports

### Before you begin

Before modifying a performance report, you require an existing report in the **Reports** list view.

To modify a performance report:

### Procedure

1. Select a storage system.
2. Select **Performance > Reports**.
3. Select a report and click **Modify**.
4. Make your changes to the report **Description**, **Symmetrix ID**, **Format**, **Generated Time Zone**, and **TimeFrame**.

You cannot change the report name.

5. Select one of the following actions in the **Queries** section:
  - Click **Add** to open the Create Query Wizard and create a new query.
  - Select a query and click **Delete**. Click **OK** in the confirmation dialog box.
  - Edit the query. Complete the following steps:

#### Note

Reports must have at least one query.

- a. Select a query and click **Edit** to open the Edit Query wizard.
- b. Make your changes in each page of the wizard.

- c. Click **OK** to close the wizard.
6. (Optional) If required, modify the schedule settings for the report.
7. If required, modify the email settings for the report.
8. Click **OK**.

For more information about reports, refer to [Creating performance reports](#) on page 579.

## Deleting performance reports

### Before you begin

You require an existing report in the Reports list view.

To delete performance reports:

### Procedure

1. Select a storage system.
2. Select **Performance > Reports**.
3. Select the report you want to delete and click **Delete**.
4. Click **OK**.

## Running performance reports

### Before you begin

You require a configured performance report in the Reports list.

To run a performance report:

### Procedure

1. Select a storage system.
2. Select **Performance > Reports**.
3. Select one or more reports and click **Run Now**.
4. Click **OK** on the **Save Report Results** dialog box.
5. 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.

---

### Note

If the report contains multiple queries involving volumes and disks the report is saved to this location:

```
/SMAS/report/SPA/runnow
```

---

## Scheduling performance reports

### Before you begin

You must have a configured performance report in the Reports list.

To run a performance report:



## Procedure

1. Select a storage system.
2. Select **Performance > Reports**.
- 3.

Select one or more reports and click  > **Schedule**.

4. Select the **Schedule** check box.
5. Configure values for each of the following options:

- **First Run Time**

The **First Run Time** is continually adjusted relative to the current time and the report execution interval. For example, if **First Run Time** is set to 3am, but the current time is 4am and the **Time Range** is set to 'Last 4 hours', then the value of **First Run Time** is changed to 7am.

- **Day(s) to run**
- **Retention**

6. Click **OK**.

Scheduled XML, CSV, and BTP reports are saved to this location:

```
SMAS\report\SPA\querydata
```

---

### Note

If the report contains multiple queries involving volumes and disks the report is saved to this location:

```
/SMAS/report/SPA/runnow
```

---

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:

### Procedure

1. Select a storage system.
2. Select **Performance > Reports**.
- 3.

Select one or more reports and click  > **Cancel Schedule**.

4. Click **OK**.

## Copying performance reports


### Before you begin

Before copying a performance report, you require an existing report in the **Reports** list view.

To copy a performance report:

### Procedure

1. Select the storage array.
2. Select **Performance** > **Reports** to open the **Reports** list view.
3.
 


 Select a report and click **> Copy**.
4. Update the report settings, as appropriate. You can edit the following data:
  - Name
  - Description
  - Symmetrix ID
  - Format
  - Generated Time Zone
  - TimeFrame
5. In the **Queries** section, do one of the following:
  - Click **Add** to open the **Create Query** wizard and create a query.
  - Select a query and click **Delete**. Click **OK** on the confirmation message.
  - Select a query and click **Edit** to open the **Edit Query** wizard.  
A report must have at least one query.
6. (Optional) In the **Schedule** tab, configure the schedule settings for the report.
7. (Optional) In the **Email** tab, configure the email settings for the report.
8. Click **OK**.

For more information about reports, see [Creating performance reports](#) on page 579

## Viewing Real Time traces

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

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:

### Procedure

1.
 


 Select  to open the **Settings** panel.

## 2. Select **Performance** > **Real Time Traces**.

Use the Real Time Traces list view to display and manage Real Time traces.

The following properties display:

### **Array**

The storage system identifier.

### **Start Time**

The start time and date 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 specific 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 saved.

### **Description**

A user-defined description of the trace.

The following controls are available:

- **Create** — [Creating a Real Time trace](#) on page 587
- **Modify** — [Modifying a Real Time Trace](#) on page 588
-  — [Deleting a Real Time trace](#) on page 588

## Creating a Real Time trace

### **Before you begin**

You must have access to one or more storage systems with Real Time registrations.

To create a Real Time trace:

### **Procedure**

1. Select a storage system.
2. Select **Performance** > **Real Time Traces**.
3. Click **Create**.
4. Select the **Symmetrix ID**.

5. 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.
6. 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**.
7. Select the **Retention** number of days.  
Values are 1–30.
8. (Optional) Enter a **Description** of this trace.
9. Click **OK**.

## Modifying a Real Time Trace

### Before you begin

You must have access to licensed storage systems.

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.

To edit a trace:

### Procedure

1. Select a storage system.
2. Select **Performance** > **Real Time Traces**.
3. Select a trace from the list and click **Modify**.
4. Update the **Trace Start Time**, **Retention** and/or **Description** values.
5. 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:

### Procedure

1. Select a storage system.
2. Select **Performance** > **Real Time Traces**.
3.  
Select a trace from the list and click  > **Delete**.
4. Click **OK**.

## Plan View

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

The Plan view displays predicted future data, based on linear projection. You can use the projection dashboards contained in the Plan view to assist when planning and to avoid potential problems.


The following projection dashboards are available from the drop-down menu:

- [Thin pool projection dashboard](#) on page 590
- [SRP projection dashboard](#) on page 589

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.

To save the current Plan view as a PDF, click  and select one of the following options:

- **Save as PDF - Charts**
- **Save as PDF - Charts & Data**

To schedule a report, click . For more information, see [Scheduling performance reports](#) on page 584.

## SRP projection dashboard

The SRP protection dashboard displays utilization data based on projections calculated over the selected time range.

### SRP - Capacity Status





This chart displays a table containing current information for all SRPs in the array. The table contains the following columns:

#### ID

The ID of the SRP.

#### Days to Full

The number of days until the capacity of the SRP 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
	< 180 days utilization remaining
	> 180 days utilization remaining

### Capacity Trend/Month

The capacity growth rate of the SRP per month. Arrow icons indicate if it's going up, down, or remaining stable.

### Capacity

The total capacity of the SRP 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 SRP. When the chart is initially opened, the first SRP is displayed, by default.

You can configure the time ranges for this chart by selecting an alternative time range. The following time ranges are available:

- 1 week

- 1 Month
- 3 Months
- 6 Months

By default, when initially opened, the **% Used Capacity Projections** chart displays data for one week.

## Thin pool projection dashboard

The thin pool protection dashboard displays utilization data based on projections calculated over the selected time range.

### Thin Pool - 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.

#### 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
	< 180 days utilization remaining
	> 180 days utilization remaining

### Capacity Trend/Month

The capacity growth rate of the thin pool per month. Arrow icons indicate if it's 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 alternative time range. The following time ranges are available:

- 1 week
- 1 Month
- 3 Months
- 6 Months

By default, when initially opened, the **% Used Capacity Projections** chart displays data for one 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.

## Viewing system registrations

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

To view system registrations:

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Performance > System Registrations** to open the Registrations list view.

Use the Registrations list view to view and manage storage system performance registration and data collection policies.

The following properties display:


- **Array** — The ID of the registered storage system.
- **Diagnostic** — A check box that indicates if the storage system is registered for Diagnostic data. The check box can be selected or cleared.
- **Real Time** — A check box that indicates if the storage system is registered for Real Time data. The check box can be selected or cleared.
- **Storage Group Real Time** — Applicable only to storage systems running PowerMaxOS 5978. A check box indicates if the storage system is registered for Storage Group Real Time data. The check box can be selected or cleared.

The overall status is captured in the All row. In this row:




- The **Diagnostic** checkbox is selected/cleared if all of the array diagnostics checkboxes are selected/cleared.
- The **Real Time** checkbox is selected/cleared if all of the array real time checkboxes are selected/cleared.
- The **Storage Group Real Time** checkbox is selected or cleared if all of the array storage group real time check boxes are selected or cleared.

---

### Note

In cases where there are a mixture of selected and cleared check boxes for items in a column, the overall status is indeterminate (denoted for the appropriate column by the  symbol).

---



3. Hover over a row and click  (see [Changing registration details](#) on page 593).
4. Hover over a row and select  (see [Viewing system registration details](#) on page 592).
5. For arrays running PowerMaxOS 5978 or later, hover over a row and select  (see [Registering storage systems](#) on page 592).

## 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:

### Procedure

1. Select .
2. Select **Performance > Settings > System Registrations**.
3. Expand the **All** list.
4. Select an array and click  to open the **Information** dialog box.

The following properties display:

- **Array** — The ID of the registered storage system.
- **Diagnostic** — Indicates if the storage system is registered for Diagnostic data .
- **Real Time** — Indicates if the storage system is registered for Real Time data.
- **Real Time Storage Groups** — Indicates if selected storage groups in the storage system are registered for real time data collection.
- **Collection Recovery (Hours)** — The amount of data that the `storstopd` daemon will save if the Performance server or the Performance database is not ready to accept new data points. The default is 24 hours.
- **Diagnostic Interval (Minutes)** — The frequency of polling in the Diagnostic view. The default 5 minutes is recommended.
- **Server IP** — The IP address of the Unisphere server.
- **SE Server IP** — The IP address of the Solutions Enabler server.

## Registering storage systems

The Performance component of Unisphere 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. The collection of real time data is performed at 5 second intervals.



If a storage system is registered for real time data collection, then by selecting a storage group, all devices in the storage group are also registered for the collection of real time data. Only one storage system can be enabled for real time storage group data collection at any given time, and a maximum of five storage groups in this array can be selected for collection.


Storage group real time data collection is available only for storage systems running PowerMaxOS 5978, or later.

To register a storage system:



**Procedure**

1. Select .
2. Select **Performance > System Registrations**.
3. Select the storage systems to register for performance data collection.
4. Select the type of performance data to collect: **Diagnostic** only, or both **Real Time** and **Diagnostic**.
5. (Optional) If **Real Time** data collection has been enabled, **Storage Group Real Time** data can also be enabled. To do this, complete the following steps:
  - a. Click .
 

The **Storage Groups for Real Time Collection** dialog box displays.
  - b. Select storage groups in the **Available Storage Groups** list and click  to move them to the **Selected Storage Groups** list.
 


A maximum of five storage groups can be selected.
  - c. Click **OK**.
 

The **Storage Group Real Time** column displays a green icon for the selected storage system.
6. Click **Apply**.

**Removing a system registration**

To remove a system registration:



**Procedure**

1. Select .
2. Select **Performance > System Registrations**.
3. Clear the registration for **Diagnostic**, **Real Time**, or both.
4. Click **Apply**.

**Viewing registered storage systems information**



To view registered storage systems information:

**Procedure**

1. Select .
2. Select **Performance > System Registrations**.
3. Hover over the storage system and click  to view the registration information.

**Changing registration details**


**Procedure**

1. Select .
2. Select **Performance > System Registrations** to open the Registrations list view.
3. Hover over the storage system name or **All**.
4. Click .
5. In the **Registration Details** dialog box, update one or both of the following fields:
  - **Collection Recovery (Hours)** — The amount of data that the `storstopd` daemon will save if the Performance server or the Performance database is not ready to accept new data points. The default is 24 hours.
  - **Diagnostic Interval (Minutes)** — The frequency of polling in the Diagnostic view. The default 5 minutes is recommended.
6. Click **OK**.

## Managing dashboard catalog

To view and edit the catalog of items displayable on the Performance dashboard:


**Procedure**

1. Select .
2. Select **Performance > Dashboard Catalog**.  
A list of categories is displayed.
3. Select the checkbox for a category to make it visible and clear a checkbox for a category to make it invisible.
4. Click **APPLY**.

## Viewing Performance databases

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

**Procedure**

1. Select  to open the **Settings** panel.
2. In the **Settings** dialog box, select **Unisphere Databases > Performance Databases**.

Use the Performance Databases list view to view and manage databases.

The following properties display:

**Array ID**

The storage system identifier.

**Status**

The status of the performance database.

**Load Status**

The status of a database action (backup, restore).

**Backup Scheduled**

Time of the scheduled backup.

The following controls are available:

- **Backup**—[Backing up a database](#) on page 597
- **Restore**—[Restoring a database](#) on page 597
- **View Details**—[Viewing database details](#) on page 595
- **Edit Scheduled Backup**—[Editing a scheduled Performance database backup](#) on page 598
- **Cancel Scheduled Backup**—[Canceling a scheduled database backup](#) on page 598
- **Edit Retention Settings**—[Editing database retention settings](#) on page 599
- **Delete**—[Deleting databases](#) on page 599
- **Remove Backup Files**—[Removing database backup files](#) on page 600


## Viewing database details

**Before you begin**

You require one or more storage system performance databases.

To view database details:

**Procedure**

1. Select  to open the **Settings** panel.
2. In the **Settings** dialog box, select **Unisphere Databases > Performance Databases**.
3. Select a database and click **View Details**.

The following details display:

**Array ID**

The storage system identifier.

**Registered to collection**

Shows the performance views available. Possible values are **Real Time** and **Diagnostic**, or **Not Registered**.

**Database administration****Last backup time**

Date and time of last backup.

**Last restore time**

Date and time of last restore.

**Last aggregation time**

Last date and time that Trending & Planning data was derived from Diagnostic data.

**Last retention time**

Date and time of last retention.

**Last real time retention time**

Date and time of last Real Time retention.

**Database size**

Size of the database.

**Historical**

**Oldest available data**

Date and time of first available Historical data.

**Latest available data**

Date and time of latest Historical data.

**Current historical data retention (months)**

Number of months to save Historical data.

**Diagnostic**

**Oldest available data**

Date and time of first available Diagnostic data.

**Latest available data**

Date and time of latest Diagnostic data.

**Current diagnostic data retention (days)**

Number of days to save Diagnostic data.

**Real Time**

**Current real time data retention (hours)**

Number hours to save real time data.

**Backup Files**

**Current backup file retention (number of days)**

Number days to save backup files.

**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 historical**

Indicates whether the last day of historical 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.

## Restoring a database

**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.

You can restore the performance database of a storage system regardless of whether it is registered.

**Procedure**


1. In the **Settings** dialog box, select **Unisphere Databases > Performance Databases**.
2. Select a storage system.
3. Click **Restore**.
4. Select a backup file.
5. Do one or more of the following:
  - Select **Restore diagnostic data if available** if you want to restore Diagnostic data.
  - Select **Restore Real Time traces if available** if you want to restore Real Time traces.
6. Click **OK**.

## 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:

**Procedure**

1. Select  to open the **Settings** panel.
2. In the **Settings** dialog box, select **Unisphere Databases > Performance Databases**.
3. Select a database from the list and click **Backup**.
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**—Complete the following steps:
  - a. Click **Next**.
  - b. Select the **Start Date** and time.
  - c. 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 **OK**.

A database backup can take some time.





## 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:

**Procedure**

1.  Select  to open the **Settings** panel.
2. In the **Settings** dialog box, select **Unisphere Databases > Performance Databases**.
3.  Select a database, click , and click **Cancel Scheduled Backup**.
4. Click **OK**.





## Editing a scheduled Performance database backup

**Before you begin**

There must be one or more database backups scheduled for a future time.

To edit a scheduled database backup:

**Procedure**



1.  Select  to open the **Settings** panel.
2. In the **Settings** dialog box, select **Unisphere Databases > Performance Databases**.
3.  Select a database, click , and click **Edit Scheduled Backup**.
4. Change the start date, time, or reoccurrence of the backup.

5. Click **OK**.

## Editing database retention settings

To edit database retention settings:

### Procedure

1. Select  to open the **Settings** panel.
2. In the **Settings** dialog box, select **Unisphere Databases > Performance Databases**.
3. Select a database, click , and click **Edit Retention Settings**.
4. In the **Retention Settings** dialog box, edit one or more of the following settings:
  - **Set historical data retention (months)**
  - **Set diagnostic data retention (days)**
  - **Set real time data retention (hours)**
  - **Set backup file retention (number of files)**
5. Click **OK**.

## Deleting databases



### Before you begin

You require one or more storage systems that are no longer registered and have a performance database.

The database delete action can be performed only 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.

To delete a performance database:

### Procedure




1. Select  to open the **Settings** panel.
2. In the **Settings** dialog box, select **Unisphere Databases > Performance Databases**.
3. Select a database from the list, click , and click **Delete**.
4. 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
5. Click **OK**.

## Removing database backup files

This procedure is for removing unneeded database backup files.

To delete database backup files:


### Procedure

1. Select  to open the **Settings** panel.
2. In the **Settings** dialog box, select **Unisphere Databases > Performance Databases**.
3.  Select a database, click , and click **Remove Backup Files**.
4. Select the file(s) to delete from the list of backup files and click **OK**.

## Viewing Performance thresholds and alerts

You can configure a warning threshold and a critical threshold value for each metric.

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Performance Thresholds and Alerts**.
3. Select All or a storage system.
4. 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.

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

The following properties display:

- **Name**—The metric name.
- **Warning**—The warning threshold.
- **Alert**—Indicates if a warning alert has been generated. The icon displayed corresponds to the alert type.
- **Critical**—The critical threshold.
- **Alert**—Indicates if a critical alert has been generated.. The icon displayed corresponds to the alert type.
- **KPI**—Indicates if the metric is a KPI.
- **Custom**—Indicates if a custom threshold or alert has been generated..

5. Click **APPLY**.

The following controls are available:

- **Create-** [Creating a performance threshold alert](#) on page 60



-  - [Deleting performance thresholds and alerts](#) on page 62.

## Performance Threshold Alert operations


- Creating a Performance Threshold Alert (see [Creating a performance threshold alert](#) on page 60).
- Editing a Performance Threshold Alert (see [Editing a performance threshold alert](#) on page 61).

## Creating a performance threshold 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 Diagnostic view.

To create performance threshold alerts:

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Performance Thresholds and Alerts**.
3. Select a storage system.
4. Select the category for which you want to create a threshold or alert.
5. Click **Create**.

The **Create Threshold and Alert** wizard displays.

6. Select the **Array**, **Category**, and **Metrics**.
- 7.

Select Instances from the **Available Instances** list and click  to move them to the **Instances to Enable** list.

8. Add a value for **Warning Threshold** or **Warning Threshold and Critical Threshold**.
9. Click **NEXT**.
10. To add an alert for each configured threshold, complete the following steps:
  - a. Select **Enable Alert**.
  - b. For each threshold you are configuring, specify values for the following fields:

### Severity

The following values are available:

- **Information**
- **Warning**
- **Critical**

### Occurrence

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.

### Samples

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.

- c. (Optional) If required, 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.



11. Click **OK**.

## Editing a performance threshold 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 threshold alerts:

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Performance Thresholds and Alerts**.
3. Navigate to the threshold alert to be edited by selecting the appropriate storage system and category.
4. Hover over an item from the table and click .
5. Edit the settings.
6. 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:


### Procedure

1. Select  to open the **Settings** panel.
2. Select **Alerts > Performance Thresholds and Alerts**.
3. Navigate to the threshold or alert to be edited by selecting the appropriate category in the **Category** section.
4. Select one or more rows and click .
5. Click **OK**.

## Configuring SNMP notifications

To set up SNMP notifications:


**Procedure**

1. Select  to open the **Settings**
2. Select **Alerts > Notifications** to open the **Notifications** page.
3. In the SNMP part of the panel, click **Configure**.
4. Click **Create**.
5. Enter an IP Address or Host Name.
6. Enter a port number.
7. To test the SNMP configuration, click **TEST**.

**Managing dashboard catalog**

To view and edit the catalog of items displayable on the Performance dashboard:

**Procedure**


1. Select .
2. Select **Performance > Dashboard Catalog**.  
A list of categories is displayed.
3. Select the checkbox for a category to make it visible and clear a checkbox for a category to make it invisible.
4. Click **APPLY**.

**Configuring email notifications**

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 notifications:

**Procedure**

1. To set up email notification:
  - a. Select  to open the **Settings** panel.
  - b. Click **Alerts > Notifications**.
  - c. In the **Email** section, click **Configure**.
  - d. In the **Outgoing Mail Server (SMTP)** section specify the following details:
    - **IP Address/Host**
    - **Server Port**
  - e. In the **User Information** section, specify the **Sender E-mail Address**.
  - f. In the **Recipients** section, click **Create** and specify the address you want to add.
  - g. Select one or more system or performance level indicators or reports to enable email notifications for the relevant level of system or performance notifications.

h. Click **APPLY**.

## About exporting and importing performance settings

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

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.

## Importing Performance settings

### 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.

---

### Note

An imported metrics definition overrides server settings. User templates are merged.

---


To import 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 page 604.

### Procedure

1. Select  to open the **Settings** panel.
2. Click **Performance > Import Settings**.
3. Select the settings to import: **Metrics and Alert Definitions, User Templates**, or both.
4. Click **Choose file** to select the settings file to import.
5. Enter the file password (which was set during the export process).

6. Click **APPLY**.

## Exporting Performance settings

### 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.



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 page 604.

To export Performance settings:

### Procedure



1.  Select  to open the **Settings** panel.
2. Click **Performance > Export Settings**.
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 **APPLY**.  
The configuration ZIP file is downloaded in your browser.
6. In the confirmation dialog box, 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:

### Procedure

1.  In the online Unisphere for PowerMax client, select  to open the **Settings** panel.
2. Click **Performance > Export PV Settings**
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:
  - **Metrics and Thresholds (default and custom)**

- **User Defined Dashboards**
- **Reports**
- **Templates**

4. Click **APPLY**.

The exported files are automatically saved in the following directories:

- **SYMAPI Database:** `install_dir\SMAS\backup\SPA\SE`
- **Settings:** `install_dir\SMAS\backup\SPA\settings`

## Metrics

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

The following table lists all metrics that are available for arrays.

**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 of 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	
RDF/A WP Count	The number of RDF/A writes pending.
Number of R2 Active Sessions	
Total Number of RDF Sessions	
Host Reads/sec	The number of host read operations performed each second by all 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.

**Table 8** Array metrics (continued)

<b>Metric</b>	<b>Description</b>
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 IOs/sec	The total IO from all BE directors to the disks per second.
BE Reqs/sec	A data transfer of a read or write between the cache and the director.
BE Read Reqs/sec	A data transfer of a read between the cache and the director.
BE Write Reqs/sec	A data transfer of a write between the cache and the director.
System WP Events/sec	The number of times each second that write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.
Device WP Events/sec	The number of times each second that the write-pending limit for a specific volume was reached.
System WP Count	The number of system cache slots that are write pending.

**Table 8** Array metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
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 MBs read each second directly from disks bypassing the cache.
Avg Optimized Read Miss Size (KB)	
FE Balance	Balance information for FE. 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.
Cache Balance	Balance information for cache. 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.
SATA Balance	Balance information for SATA. 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 (continued)

<b>Metric</b>	<b>Description</b>
FC Balance	Balance information for FC. 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.
EFD Balance	Balance information for EFD. 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.
DX Balance	Balance information for DX. 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.
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	Balance information for DA. 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.
Allocated Capacity	The total allocated array capacity.
Critical Alert Count	Number of critical alerts.
Warning Alert Count	Number of warning alerts.
Information Alert Count	Number of information alerts.
Disk Utilization	Percentage of time the disk is busy.
RDF Utilization	Percentage of time RDF is busy.
BE Utilization	Percentage of time the BE is busy.
FE Utilization	Percentage of time the FE is busy.
Copy Slot Count	
Overall Efficiency Ratio	Ratio of the sum of all TDEVs + snapshot sizes (calculated based on the 128K track size) and the Physical Used Storage (calculated based on the compressed pool track size).
RDF Balance	Balance information for RDF. Balance metrics measure hardware components balance. It is calculated as the difference between the maximum and minimum values

**Table 8** Array metrics (continued)

Metric	Description
	for "% Busy" for the component instances. In a balanced system the balance metrics should be relatively low.
% Snapshot Saved	Percentage savings of the sum of all TDEV Snapshot sizes (at the time of Snapshot creation) and the TDEV Snapshot Allocated Storage.
% Virtual Provisioning Saved	Percentage savings of the TDEV configured storage presented to the hosts and the TDEV Allocated Storage.
Snapshot Compression Ratio	<p>Ratio of the RDP Logical Backend Storage (calculated based on the 128K track size) and the RDP Physical Used Storage of the RDP space (calculated based on the compressed pool track size).</p> <hr/> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p> <hr/>
Snapshot Efficiency Ratio	Ratio of the sum of all snapshot sizes and the RDP Logical Backend Storage.
Snapshot Shared Ratio	Ratio of the Snapshot Allocated Storage and the RDP Logical Backend Storage. Sum of all modified tracks on all snapshots / Sum of the RDP size for all snapshot.
Virtual Provisioning Compression Ratio	<p>Ratio of the TDEV Logical Backend Storage (calculated based on the 128K track size) and the TDEV Physical Used Storage (calculated based on the compressed pool track size).</p> <hr/> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p> <hr/>
Virtual Provisioning Efficiency Ratio	Ratio of the TDEV configured storage and the TDEV Logical Backend Storage (calculated based on TDEV allocated tracks without shared unowned).
Virtual Provisioning Shared Ratio	Ratio of the TDEV allocated storage and the TDEV Logical Backend Storage. Backend Storage is the sum of TDEV allocated tracks without counting the TDEV shared unowned tracks.
Overall Compression Ratio	Ratio of the sum of all TDEVs + snapshot sizes (calculated based on the 128K track size) and the Physical Used Storage (calculated based on the compressed pool track size).

**Table 8** Array metrics (continued)

Metric	Description
	<p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p>
% Subscribed Capacity	The subscribed capacity as a percentage of the usable capacity.
% Effective Used Capacity	The percentage of usable capacity that would be used if all compressible data was compressed.
Usable Capacity GB	The total capacity of all Storage Resource Pools (SRP) in GB.
% Metadata Replication Used	The percentage of metadata used for replication.
% Metadata System Used	The percentage of metadata used for the system.
Glacial Track Count	The number of tracks currently in glacially compressed state.
% Meta Data Backend Used	BE TID used percent.
% Meta Data Frontend Used	FE TID used percent.
Snapshot Capacity (GB) + (TB)	Total snapshot capacity for system data.
Snapshot Modified Capacity (GB) + (TB)	Total snapshot modified capacity for system data.
Used Usable Capacity (GB) + (TB)	Total user used capacity.
Capacity Health Score	Capacity Health Score.
Hardware Config Health Score	Hardware Config Health Score.

## 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

**Table 9** BE director (DA) metrics (continued)

Metric	Description
	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.
Clone Copy Read	
Clone Copy Write	
PHCO Rebuild Copy	Internal metric.
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.

**Table 9** BE director (DA) metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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 Work Time	
% Busy	The percent of time the director 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	The number of read operations performed each second by the port.
Writes/sec	The number of write operations performed each second by the port.
Host IOs/sec	The number of host operations performed each second by the port.
MBs Read/sec	The read throughput (MBs) of the port per second.
MBs Written/sec	The write throughput (MBs) of the port per second.
Host MBs/sec	The size of the data transfer from the host in MBs per second.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Speed	Total port capacity. Note that in the REST API, for backward compatibility reasons, this metric is named <b>Max Speed Gb/sec</b> .
% Busy	The percent of time that the port is 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 maximum 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	Outbound speed.

**Table 12** Board metrics (continued)

Metric	Description
Inbound MBs/sec	Inbound speed.
Max Speed MBs/sec	Maximum speed.
Utilization	Calculated value: $((\text{transfer-rate-in} + \text{transfer-rate-out}) / \text{max-transfer-rate})$

## Cache Partition metrics

The following table lists all metrics that are available for cache partitions.

**Table 13** Cache partition metrics

Metric	Description
Acquired Slot Count	
Age Non WP Slots	
Avg Age Given Destage	
Avg Age of Write to Non WP Slot	$(\text{age non wp slots}) / ((\text{writes to all non wp slots}) - (\text{writes to young non wp slots}))$
Cache Age GT 10 Min	
Cache Age GT 1 hour	
Cache Age GT 1 Min	
Cache Hit Ages 1	
Cache Hit Ages 2	
Cache Hit Ages 3	
Cache Hit Ages 4	
Cache Hit Ages 5	
Cache Hit Ages 6	
Cache Hit Ages 7	
Cache Hit Ages 8	
Cache Slot Ages 1	
Cache Slot Ages 2	
Cache Slot Ages 3	
Cache Slot Ages 4	
Cache Slot Ages 5	
Cache Slot Ages 6	
Cache Slot Ages 7	
Cache Slot Ages 8	

**Table 13** Cache partition metrics (continued)

Metric	Description
Cache Slots Used	
Destage Slot Age	
Destaged Slot Count	
Donation Give Count	
Donation Take Count	
Donation Time	
Host IOs/sec	Host operations performed each second by the cache partition.
Local WP Count	The number of write pending slots waiting to be de-staged to disk on the local system. The value should be less than or equal to the system write pending count.
% Max Cache	Maximum slot allocation for a partition.
Host MBs/sec	The number of host MBs written and read each second.
% Min Cache	Minimum slot allocation for a partition
% 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 percent of IO operations that were immediately satisfied from cache.
Remote WP count	The number of write pending slots waiting to be de-staged to disk on the remote system. The value should be less than or equal to the system write pending count.
% Cache Used	The percent of the cache partition that is used.
% Target Cache	Percent of total cache allocated to this partition.
Total Replace Slots	
WP Count	The number of cache partition slots that are write pending.
WP Limit	The maximum number of write-pending slots.
Writes to All Non WP Slots	
Writes to Young Non WP Slots	

## DATA Volume metrics

The following table lists all metrics that are available for DATA volumes.



**Table 14** DATA volume metrics

<b>Metric</b>	<b>Description</b>
Allocated Capacity (GB)	The total allocated volume capacity.
BE MBs Read/sec	The number of MBs read by the disk directors from the volume each second.
BE MBs Transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)
BE MBs Written/sec	The number of MBs written to the volume from the disk director each second.
BE % Reads	The percent of the back-end IO that were read requests.
BE % Writes	The percent of the back-end IO that were write requests.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the 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 the cache.
Kbytes Written for IVTOC	
BE Reads For Copy (KB)	Internal metric.
Kbytes Read for Rebuild	
BE Reads For VLUN Migration (KB)	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Writes For Rebuild (KB)	Internal metric.
BE Writes For VLUN Migration (KB)	Internal metric.
BE Optimize Writes (KB)	Internal metric.
BE Partial Sector Writes (KB)	Internal metric.
Reads for Rebuild	
BE Reads For VLUN Migration	Internal metric.
BE Reads For Copy	Internal metric.
BE Writes For Copy	Internal metric.
BE Writes For Rebuild	Internal metric.
BE Writes For VLUN Migration	Internal metric.
BE XOR Reads (KB)	Internal metric.
BE XOR Reads	Internal metric.
Device Block Size	The block size of the volume.
Total Capacity (GB)	The total capacity of the volume (GBs).

**Table 14** DATA volume metrics (continued)

Metric	Description
IO Density	The number of BE requests per GB of disk.
Parity Generation Kbytes Read	
Parity Generation Reads	
BE Prefetched MBs/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.
BE Reqs/sec	The number of read/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.
Writes for IVTOC	

## Database metrics

The following table lists all metrics that are available for databases.

**Table 15** 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.
Writes for IVTOC	
RDF Writes/sec	Total RDF writes per second for the database.
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.

**Table 15** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
Kbytes written for IVTOC	
Read RT (ms)	The average time that it took the system to serve one read command.
Write RT (ms)	The average time that it took the system to serve one write command.
Read Miss RT (ms)	The average time a read miss operation was performed by the group. 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 group. 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	Percentage of IO operations that were reads.
% Writes	Percentage of IO operations that were writes.
% Read Hit	The percentage of read operations performed that were immediately satisfied by cache.
% Write Hit	The percentage of write operations performed that were immediately satisfied by cache.
% Read Miss	The percentage of read miss operations performed, a miss occurs when the requested read data is not found in cache or the write

**Table 15** Database metrics (continued)

Metric	Description
	operation had to wait while data was destaged from cache to the disks.
% Write Miss	The percentage of write miss operations performed, 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.
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 that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed each that were misses.
Seq Write Hits/sec	The number of sequential write operations performed by the volume that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by 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.

**Table 15** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.

**Table 15** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
Parity Generation Kbytes Read	
Parity Generation Reads	
% 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.
% Random Reads	The percent of read commands from a host not identified as part of a sequential stream.
% Random Writes	The percent of write commands from a host not identified as part of a sequential stream.
BE % Reads	The percent of read operations from the back-end directors that were immediately satisfied by cache.
BE % Writes	The percent of write operations from the back-end directors that were immediately satisfied by cache.
Write RT Count 5 to 8	
Write RT Count 8 to 14	
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	

**Table 15** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
Response Time (ms)	The average response time for the reads and writes.
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	Number of information alerts.
Critical Alert Count	Number of critical alerts.
Warning Alert Count	Number of information alerts.
Allocated Capacity (GB)	The total allocated database capacity.
BE Partial Sector Writes (KB)	The KBs of back-end IO that were partial sector writes.
IO Density	The number of BE requests per GB of disk.
BE 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 Reads For Rebuild (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.

**Table 15** Database metrics (continued)

Metric	Description
BE RDF Copy (MB)	Internal metric.
Reads For Rebuild	
Total Capacity (GB)	The total capacity of the database.
Device Block Size	The block size of the volume.
BE Prefetched MBs/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.
RDF Read Hits/sec	
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.
RDF MBs Read/sec	
RDF Rewrites/sec	
RDF MBs Written/sec	
Avg Optimized Read Miss Size (KB)	
Write Paced Delay (ms)	

## Database by Pool metrics

The following table lists all metrics that are available for databases by pool.

**Table 16** Database 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 allocated pool 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	The number of writes per second that were destaged to disk.



**Table 16** Database by pool metrics (continued)

Metric	Description
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.
Destage Write MBs/sec	The size (MBs) of writes per second that were destaged to disk.
Prefetched Tracks/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.
Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
Total Tracks	The number of allocated tracks in the database by pool.
Compressed Tracks	The number of compressed tracks in the database by pool (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the total tracks that are compressed (applies only to storage systems running Enginuity OS 5876).

## Device Group metrics

The following table lists all metrics that are available for device groups.

**Table 17** 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	Total RDF writes per second for the device group.
Host Writes/sec	The number of host write operations performed each second by the group.
Host Hits/sec	The number of host read/write operations performed each second by the group that were immediately satisfied from cache.
Host Read Hits/sec	The number of host read operations performed each second by the group that were immediately satisfied from cache.

**Table 17** Device group metrics (continued)

<b>Metric</b>	<b>Description</b>
Host Write Hits/sec	The number of host write operations performed each second by the group that were immediately satisfied from cache.
Host Misses/sec	The number of host read/write operations performed each second by the group that could not be satisfied from cache.
Host Read Misses/sec	The number of host read operations performed each second by the group that were not satisfied from cache.
Host Write Misses/sec	The number of host write operations performed each second by the group that were not satisfied from cache.
Host MBs/sec	The cumulative number of host MBs read/writes per second by the group.
Host MBs Read/sec	The cumulative number of host MBs read per second by the group.
Host MBs Written/sec	The cumulative number of host MBs written per second by the group.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read 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.

**Table 17** Device group metrics (continued)

<b>Metric</b>	<b>Description</b>
% Write Hit	The percent of write operations, performed by the group, that were immediately satisfied by cache.
% Read Miss	The percent of read miss operations performed each second by the group. A miss occurs when the requested read data is not found in cache.
% Write Miss	The percent of write miss operations performed each second by the group. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.
WP Count	The number of tracks currently in write pending mode for the group.
Seq IOs/sec	Number of IO operations performed each second that were sequential.
Seq Reads/sec	Number of read IO operations performed each second that were sequential.
Seq Writes/sec	Number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	Number of sequential read operations performed each second by the group that were immediately satisfied from cache.
Seq Read Miss/sec	Number of sequential read operations performed each second by the group that were misses.
Seq Write Hits/sec	Number of sequential write operations performed each second by the group that were immediately satisfied from cache.
Seq Write Misses/sec	Number of sequential write operations performed each second by the group that were misses.
Random IOs/sec	IOs from a host not identified as part of a sequential stream.
Random Reads/sec	Read IO commands from a host not identified as part of a sequential stream.
Random 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.

**Table 17** Device group metrics (continued)

<b>Metric</b>	<b>Description</b>
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	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
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.

**Table 17** Device group metrics (continued)

<b>Metric</b>	<b>Description</b>
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 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

**Table 17** Device group metrics (continued)

Metric	Description
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)	The total capacity of the device group.
Allocated Capacity (GB)	The allocated capacity of the device group.
Device Block Size	The block size of the volume.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
RDF Rewrites/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Read Hits/sec	
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	

## Disk metrics

The following table lists all metrics that are available for disks.

**Table 18** 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.

**Table 18** Disk metrics (continued)

Metric	Description
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.
Avg Hypers per Seek	The average number of hypervolumes that the disk head crossed during one seek.
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 19** 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

**Table 19** Disk bucket metrics (continued)

Metric	Description
	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 20** Disk group metrics

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)
IO Density	The number of BE requests per GB of disk.
IOs/sec	The number of IO commands to the disk.
MBs/sec	The total IO (reads and writes) per second in MBs.
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.
% Busy	The percent of time that the disk is busy serving IOs.
% Idle	The percent of time the disk is idle.



**Table 20** Disk group metrics (continued)

Metric	Description
% Free Capacity	The percent of the disk that is free.
% Used Capacity	The percent of the disk group capacity that is allocated.
Reads/sec	The number of host reads per second for the disk.
Writes/sec	The number of host writes per second for the disk.
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.
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 each second.
Total Capacity (GB)	The total capacity of the disk (GBs).
Used Capacity (GB)	The total used capacity of the disk (GBs).

## Disk Group tier metrics

The following table lists all metrics that are available for disk group tiers.

**Table 21** 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.

**Table 21** Disk group tier metrics (continued)

Metric	Description
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.
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 + 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 22** 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.

**Table 22** Disk technology metrics (continued)

<b>Metric</b>	<b>Description</b>
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 + 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 23** 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 24** DX emulation metrics

Metric	Description
Total Work Time	
% Busy	The percent of time the directory is busy.

## DX Port metrics

The following table lists all metrics that are available for DX ports.

**Table 25** DX 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	
% Busy	

## EDS Director metrics

The following table lists all metrics that are available for EDS directors.

**Table 26** EDS director metrics

Metric	Description
Random Read Misses	The number of random read IOs that were misses.
Random Read Misses Mbytes	The number of random read IOs that were misses in MBs.
Random Write Misses	The number of random write IOs that were misses.
Random Write Misses Mbytes	The number of random write IOs that were misses in MBs.
% Busy	The percent of time the director is busy.
Glacial Tracks In	Ingress to glacial tier.
Glacial Tracks Out	Egress from glacial tier.

## EDS Emulation metrics

The following table lists all metrics that are available for EDS emulations.

**Table 27** EDS director metrics

Metric	Description
Total Work Time	
% Busy	The percent of time the directory is busy.

## External Director metrics

The following table lists all metrics that are available for external directors.

**Table 28** External director 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.
Clone Copy Read	
Clone Copy Write	

**Table 28** External director metrics (continued)

<b>Metric</b>	<b>Description</b>
PHCO Rebuild Copy	Internal metric.
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.

**Table 28** External director metrics (continued)

Metric	Description
% Non-IO Busy	The percent of time the director was busy with non-IO tasks, such as FAST swaps, Optimizer moves, Snap copies, etc.

## External Disk metrics

The following table lists all metrics that are available for external disks.

**Table 29** 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.



**Table 29** External disk metrics (continued)

<b>Metric</b>	<b>Description</b>
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 30** External disk group metrics

<b>Metric</b>	<b>Description</b>
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.

**Table 30** External disk group metrics (continued)

Metric	Description
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)	The average response time for the reads and writes.
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.
Members	The number of members in the external disk group.

## FAST VP Policy metrics

The following table lists all metrics that are available for FAST VP policies.

**Table 31** 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.

**Table 31** FAST VP policy metrics (continued)

Metric	Description
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 32** FE director metrics

Metric	Description
% Busy	The percent of time the director 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).

**Table 32** FE director metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
% 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

**Table 32** FE director metrics (continued)

Metric	Description
	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	The percent of time that core 0 is busy serving IOs.
Syscall Time/sec	Total time spent processing all system calls.
Avg RDF Write Time (ms)	
Avg WP Disconnect Time (ms)	
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	The percent of time that core 1 is busy serving IOs.
% 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 1toRead RT Over 64	These buckets show the distribution of the total average read response time.
Write RT 0 to 1toWrite RT Over 64	These buckets show the distribution of the total average write response time.

**Table 32** FE director metrics (continued)

Metric	Description
Host IO Limit IOs/sec	The number of IOs the FE director is processing.
Host IO Limit MBs/sec	The number of MBs the FE port is processing.
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: $Sx=1..9 [(Avg Queue Depth Range X - 1) / Avg Queue DepthRange X] * (Queue Depth Count Range X / Queue Depth Count Total) * 100$
CCWS Per IO	
Avg Optimized Read Miss Size (KB)	
Optimized MBs Read Miss/sec	Number of host MBs read each second directly from disks bypassing the cache.
Optimized Read Misses/sec	Number of read requests each second performed directly from disks bypassing the cache.
ORS ASync IOs/sec	
ORS ASync MBs/sec	
ORS Sync MBs/sec	
ZHPF reads	
ZHPF Reads MBs	
ZHPF Writes	
ZHPF Writes MBs	

**Table 33** FE director metrics

Metric	Description
% Busy	The percent of time the director 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

**Table 33** FE director metrics (continued)

Metric	Description
	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.
% 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.

**Table 33** FE director metrics (continued)

<b>Metric</b>	<b>Description</b>
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	The percent of time that core 0 is busy serving IOs.
Syscall Time/sec	Total time spent processing all system calls.
Avg RDF Write Time (ms)	
Avg WP Disconnect Time (ms)	
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	The percent of time that core 1 is busy serving IOs.
% Write Req Hit	The percent of write requests that were satisfied from cache.



**Table 33** FE director metrics (continued)

Metric	Description
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 1toRead RT Over 64	These buckets show the distribution of the total average read response time.
Write RT 0 to 1toWrite RT Over 64	These buckets show the distribution of the total average write response time.
Host IO Limit IOs/sec	The number of IOs the FE director is processing.
Host IO Limit MBs/sec	The number of MBs the FE port is processing.
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: $S_x = 1..9 [(Avg Queue Depth Range X - 1) / Avg Queue Depth Range X] * (Queue Depth Count Range X / Queue Depth Count Total) * 100$
CCWS Per IO	
Avg Optimized Read Miss Size (KB)	
Optimized MBs Read Miss/sec	Number of host MBs read each second directly from disks bypassing the cache.
Optimized Read Misses/sec	Number of read requests each second performed directly from disks bypassing the cache.
ORS ASync IOs/sec	
ORS ASync MBs/sec	
ORS Sync MBs/sec	
ZHPF reads	
ZHPF Reads MBs	
ZHPF Writes	
ZHPF Writes MBs	
EPL Busy	

**Table 33** FE director metrics (continued)

Metric	Description
SPA Busy	

## FE Emulation metrics

The following table lists all metrics that are available for FE emulations.

**Table 34** FE emulation metrics

Metric	Description
Total Work Time	
% Busy	The percent of time the directory is 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	The average number of host reads performed per second.
Writes/sec	The average number of host writes performed per second.
Host IOs/sec	The number of IOs the port is performing on behalf of the specific SG.
MBs Read/sec	The throughput in MBs read per second.
MBs Written/sec	The throughput in MBs written per second.
Host MBs/sec	The number of MBs the port is processing on behalf of the specific SG.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Speed Gb/sec	
Max Speed Gb/sec	
% Busy	The percent of time the port is busy.
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.
Response Time (ms)	The average response time for the reads and writes.

**Table 35** FE port metrics (continued)

Metric	Description
iSCSI Checksum Error Count	The percent of time that the port is busy.

## 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	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	

## 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	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	

## FICON Emulation metrics

The following table lists all metrics that are available for FICON emulations.

**Table 38** FICON emulation metrics

Metric	Description
% Busy	The percent of time the FICON channel is busy.

## FICON Emulation Threads metrics

The following table lists all metrics that are available for FICON emulation threads.

**Table 39** FICON emulation threads metrics

Metric	Description
% Busy	The percent of time the FICON emulation threads were busy.
% Idle	The percent of time the FICON emulation threads were idle.

## FICON Port Threads metrics

The following table lists all metrics that are available for FICON emulations.

**Table 40** FICON port threads metrics

Metric	Description
% Busy	The percent of time the FICON port was busy.
% Idle	The percent of time the FICON port was idle.

## Host metrics

The following table lists all metrics that are available for hosts.

**Table 41** Host metrics

Metric	Description
Host IOs/sec	The total number of host read IO and write IO operations performed each second by the host.
Host Reads/sec	The total number of host read IO operations performed each second by the host.
Host Writes/sec	The total number of host write IO operations performed each second by the host.
Host MBs/sec	The size of the data transfer from the host in MBs per second.
Read RT (ms)	The average time it took the system to serve one read IO for this host.
Write RT (ms)	The average time it took the system to serve one write IO for this host.
Syscall Count/sec	The total number of calls seen by this host.
Host MBs Read/sec	Cumulative number of host MBs read per second.

**Table 41** Host metrics (continued)

Metric	Description
Host MBs Written/sec	Cumulative number of host MBs written per second.
Response Time (ms)	The average response time for the reads and writes.

## Host IO Limit by FE metrics

The following table lists all metrics that are available for host IO limits (by FE).

**Table 42** 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 43** 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 44** IM director metrics

Metric	Description
% Busy	The percent of time the director is busy.

**Table 44** IM director metrics (continued)

Metric	Description
Memory Page Usage	
Memory Page Fall thru Time	

## IM Emulation metrics

The following table lists all metrics that are available for IM emulations.

**Table 45** IM emulations metrics

Metric	Description
Total Work Time	
% Busy	The percent of time the directory is busy.

## Initiator metrics

The following table lists all metrics that are available for initiators.

**Table 46** Initiator metrics

Metrics	Description
Host IOs/sec	A host command for data transfer.
Host Reads/sec	The number of host read operations performed each second by the initiator.
Host Writes/sec	The number of host write operations performed each second by the initiator.
Host MBs/sec	A cumulative number of host MBs read/writes per second.
Read RT (ms)	The average time it took the system to serve one read IO for this initiator.
Write RT (ms)	The average time it took the system to serve one write IO for this initiator.
Syscall Count/sec	The total number of calls seen by this initiator.
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.
Response Time (ms)	The average response time for the reads and writes.

## Initiators by Port metrics

The following table lists all metrics that are available for initiators (by port).

**Table 47** Initiators (by port) metrics

Metrics	Description
Host IOs/sec	A host command for data transfer.
Host Reads/sec	The number of host read operations performed each second by the initiator (by port).
Host Writes/sec	The number of host write operations performed each second by the initiator (by port).
Host MBs/sec	A cumulative number of host MBs read/writes per second.
Read RT (ms)	The average time it took the system to serve one read IO for this initiator (by port).
Write RT (ms)	The average time it took the system to serve one write IO for this initiator (by port).
Syscall Count/sec	The total number of calls seen by this initiator (by port).
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.
Response Time (ms)	The average response time for the reads and writes.

## IP Interface metrics

The following table lists all metrics that are available for IP interfaces.

**Table 48** IP interface metrics

Metric	Description
Packets/sec	
MBs/sec	
Duplicate Acks Received/sec	
TCP Retransmits/sec	
Reads/sec	The number of read operations performed each second.
Writes/sec	The number of write operations performed each second.
Host IOs/sec	The number of host operations performed each second.
Total Read Time	Total time spent performing reads.

**Table 48** IP interface metrics (continued)

Metric	Description
Total Write Time	Total time spent performing writes.
Response Time (ms)	The average response time for the reads and writes.

## iSCSI Target metrics

The following table lists all metrics that are available for iSCSI targets.

**Table 49** iSCSI target metrics

Metric	Description
Packets/sec	
MBs/sec	
Duplicate Acks Received/sec	
TCP Retransmits/sec	
Reads/sec	The number of read operations performed each second.
Writes/sec	The number of write operations performed each second.
Host IOs/sec	The number of host operations performed each second.
Total Read Time	Total time spent performing reads.
Total Write Time	Total time spent performing writes.
Response Time (ms)	The average response time for the reads and writes.

## Masking View metrics

The following table lists all metrics that are available for masking views.

**Table 50** Masking view metrics

Metric	Description
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.
Response Time (ms)	The average response time for the reads and writes.
Read RT (ms)	The average time it takes to serve one read IO.



**Table 50** Masking view metrics (continued)

Metric	Description
Write RT (ms)	The average time it takes to serve one write IO.
Capacity (GB)	The total capacity of the masking group in GBs.

## Metas metrics

The following table lists all metrics that are available for metas.

**Table 51** 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.

**Table 51** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
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.

**Table 51** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
% Write Miss	The percent of write miss operations performed each second by the Symmetrix volume. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.
WP Count	The number of tracks currently in write pending mode for the volume.
Seq IOs/sec	The number of IO operations performed each second that were sequential.
Seq Reads/sec	The number of read IO operations performed each second that were sequential.
Seq Writes/sec	The number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	The number of sequential read operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed each second by the Symmetrix volume that were misses.
Seq Write Hits/sec	The number of sequential write operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by the Symmetrix volume that were misses.
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.

**Table 51** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the Symmetrix volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.

**Table 51** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
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	The number of BE requests per GB of disk.
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.

**Table 51** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	The allocated capacity of the meta.
Total Capacity (GB)	The total capacity of the meta.
Device Block Size	The block size of the volume.
XtremSW Cache Read Hits/sec	The cumulative number of reads that were hits by XtremSW Cache per second.
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by XtremSW Cache per second.
XtremSW Cache Reads/sec	The average time it took XtremSW Cache to serve one read.
XtremSW Cache Writes/sec	The average time it took XtremSW Cache to serve one write.
XtremSW Cache IOs/sec	The average time it took XtremSW Cache to serve one IO.
XtremSW Cache Skipped IOs/sec	The number of IOs that were skipped by XtremSW Cache per second.
XtremSW Cache Dedup Hits/sec	The number of XtremSW Cache Dedup hits per second.
XtremSW Cache Dedup Reads/sec	The number of XtremSW Cache Dedup reads per second.
XtremSW Cache Dedup Writes/sec	The number of XtremSW Cache Dedup writes per second.
% XtremSW Cache Reads	The percent of XtremSW Cache IO that were reads.

**Table 51** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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	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 (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 52** 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 performed by the disk directors to the cache.
Parity Generation Reads	



**Table 52** Pool-bound volumes metrics (continued)

<b>Metric</b>	<b>Description</b>
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 mode for the volume.
Seq IOs/sec	The number of IO operations performed each second that were sequential.

**Table 52** Pool-bound volumes metrics (continued)

<b>Metric</b>	<b>Description</b>
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)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$

**Table 52** Pool-bound volumes metrics (continued)

<b>Metric</b>	<b>Description</b>
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the Symmetrix volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes read requests from the disk directors to cache.
BE Disk Read 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.

**Table 52** Pool-bound volumes metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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	The number of BE requests per GB of disk.
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	

**Table 52** Pool-bound volumes metrics (continued)

Metric	Description
BE Writes For VLUN Migration (KB)	Internal metric.
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)	The allocated pool capacity in GBs.
Total Capacity (GB)	The total capacity of the pool in GBs.
Device Block Size	The block size of the volume.
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 MBs read each second directly from disks bypassing the cache.
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 53** Pool by storage group metrics

Metric	Description
BE Reqs/sec	The number of data transfers between cache and the pool by storage group.
Allocated Capacity	The allocated capacity for the pool by storage group in GBs.
IO Density	The number of BE requests per GB of disk. (BE Reads + destaged writes) / capacity

**Table 53** Pool by storage group metrics (continued)

Metric	Description
BE Read Reqs/sec	A data transfer of a read between the cache and the director.
Destage Writes/sec	The number of writes per second that were destaged to disk.
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.
Destage Write MBs/sec	The size (MBs) of writes per second that were destaged to disk.
Prefetched Tracks/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.
Pretetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
Allocated Tracks	The number of allocated tracks in the pool.
Compressed Tracks	The number of compressed tracks (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the total tracks that are compressed (applies only to storage systems running Enginuity OS 5876).
Local Replication Tracks	

## Port Group metrics

The following table lists all metrics that are available for port groups.

**Table 54** Port group metrics

Metric	Description
Reads/sec	The average number of host reads performed per second.
Writes/sec	The average number of host writes performed per second.
Host IOs/sec	The number of host IO operations performed each second, including writes and random and sequential reads.
MBs Read/sec	The number of reads per second in MBs.
MBs Written/sec	The number of writes per second in MBs.
Host MBs/sec	The number of host MBs read each second.

**Table 54** Port group metrics (continued)

Metric	Description
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
% Busy	The percent of time the port group is busy.

## RDF Director metrics

The following table lists all metrics that are available for RDF directors.

**Table 55** 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.
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	
Tracks Received/sec	The number of tracks received by this director per second.

**Table 55** RDF director metrics (continued)

<b>Metric</b>	<b>Description</b>
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 (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 MBs Sent/sec	The number of async MBs sent per second via this RDF director.
% Compressed MBs Sent and Received/sec	Calculated value: ((Compressed MBs Sent and Received/sec / MBs Sent and Received/sec) * 100)
% Compressed MBs Received/sec	Calculated value: ((Compressed MBs Received/sec / MBs Received/sec) * 100)
% Compressed MBs Sent/sec	Calculated value: ((Compressed MBs Sent/sec / MBs Sent/sec) * 100)
Compressed MBs Sent and Received/sec	The number of compressed MBs sent and received by this RDF director.
Compressed MBs Received/sec	The number of compressed MBs/sec received by this RDF director.
Compressed MBs Sent/sec	The number of compressed MBs/sec sent by this RDF director.
Number of Compressed Links	The number of compressed links used by this RDF director.
Number of Links	The number of links used by this RDF director.
Reqs/sec	The average number of requests performed by the director per second.



## RDF Emulation metrics

The following table lists all metrics that are available for RDF emulations.

**Table 56** RDF emulation metrics

Metric	Description
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 57** RDF port metrics

Metric	Description
Reads/sec	The number of read operations performed by the port per second.
Writes/sec	The number of write operations performed by the port per second.
Host IOs/sec	The number of host operations performed each second by the port.
MBs Read/sec	The read throughput (MBs) of the port per second.
MBs Written/sec	The write throughput (MBs) of the port per second.
Host MBs/sec	The size of the data transfer from the host in MBs per second.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Speed Gb/sec	
Max Speed Gb/sec	
% Busy	The percent of time that the port is busy.

## RDF/A Group metrics

The following table lists all metrics that are available for RDF/A groups.

**Table 58** 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.)

**Table 58** RDF/A group metrics (continued)

<b>Metric</b>	<b>Description</b>
RDF Reads/sec	
Avg Cycle Size	The average size of all the cycles in RDFA sessions.
Avg IO Service Time (ms)	The average time the director takes to serve IO.
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.
System WP Limit	The maximum number of write-pending slots for the entire storage system.
System WP Count	The number of system cache slots that are write pending.
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.
DSE Used Tracks	The active cycle spilled count plus the inactive cycle spilled count.

**Table 58** RDF/A group metrics (continued)

<b>Metric</b>	<b>Description</b>
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).
RDF/A Session Indicator	Indicates if group is RDFA (0/1).
RDF/A Active Indicator	Indicates if RDF/A group is active (0/1).
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.
R1 to R2 Lag Time	The time that R2 is behind R1 (RPO). This is calculated as the last cycle time plus the time since last switch. In a regular case, the cycles switch every ~30 sec and the samples are taken every few minutes, therefore this counter may not show very significant data, however, when cycles elongate beyond the sample time, this counter can help indicate an estimate of the RPO.
Num Devs in Session	The number of 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

**Table 58** RDF/A group metrics (continued)

Metric	Description
	one of the volumes reaches its WP limit. The importance of this counter is that it shows one of the common reasons for an increase in the cycle time.
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.
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 number 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)
RDF Rewrites/sec	
RDF MBs Written/sec	
RDF MBs Read/sec	

**Table 58** RDF/A group metrics (continued)

Metric	Description
RDF Read Hits/sec	
Total Tracks Spilled Back	
Total Tracks Spilled Over	
Total Tracks Spilled Over Used	

## RDF/S Group metrics

The following table lists all metrics that are available for RDF/S groups.

**Table 59** RDF/S group metrics

Metric	Description
Host IOs/sec	The number of host operations performed each second by the RDF/S Group.
Host Reads/sec	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.

**Table 59** RDF/S group metrics (continued)

<b>Metric</b>	<b>Description</b>
Host MBs Read/sec	The cumulative number of host MBs read per second by the group.
Host MBs Written/sec	The cumulative number of host MBs written per second by the group.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read 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.

**Table 59** RDF/S group metrics (continued)

<b>Metric</b>	<b>Description</b>
Seq IOs/sec	Number of IO operations performed each second that were sequential.
Seq Reads/sec	Number of read IO operations performed each second that were sequential.
Seq Writes/sec	Number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	Number of sequential read operations performed each second by the group that were immediately satisfied from cache.
Seq Read Miss/sec	Number of sequential read operations performed each second by the group that were misses.
Seq Write Hits/sec	Number of sequential write operations performed each second by the group that were immediately satisfied from cache.
Seq Write Misses/sec	Number of sequential write operations performed each second by the group that were misses.
Random IOs/sec	IOs from a host not identified as part of a sequential stream.
Random Reads/sec	Read IO commands from a host not identified as part of a sequential stream.
Random 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	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$

**Table 59** RDF/S group metrics (continued)

<b>Metric</b>	<b>Description</b>
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
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.



**Table 59** RDF/S group metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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.
RDF Response Time (ms)	
Response Time (ms)	The average time it takes to satisfy IO requests.
IO Density	The number of partial sector writes by the disk directors.
BE Partial Sector Writes (KB)	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)	The total capacity of the device group.
Device Block Size	The block size of the volume.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
RDF MBs Read/sec	
RDF Read Hits/sec	

**Table 59** RDF/S group metrics (continued)

Metric	Description
RDF MBs Written/sec	
RDF Rewrites/sec	
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 60** 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.

**Table 60** SAVE volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
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

**Table 60** SAVE volume metrics (continued)

<b>Metric</b>	<b>Description</b>
	operation had to wait while data was destaged from cache to the disks.
% Write Miss	The percent of write miss operations performed each second by the Symmetrix volume. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.
WP Count	The number of tracks currently in write pending mode for the volume.
Seq IOs/sec	The number of IO operations performed each second that were sequential.
Seq Reads/sec	The number of read IO operations performed each second that were sequential.
Seq Writes/sec	The number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	The number of sequential read operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed each second by the Symmetrix volume that were misses.
Seq Write Hits/sec	The number of sequential write operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by the Symmetrix volume that were misses.
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.

**Table 60** SAVE volume metrics (continued)

<b>Metric</b>	<b>Description</b>
Random Write Misses/sec	The number of random write IOs that were misses.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the Symmetrix volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.

**Table 60** SAVE volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	
IO Density	The number of BE requests per GB of disk.
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.

**Table 60** SAVE volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	Internal metric.
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)	The allocated volume capacity in GBs.
Total Capacity (GB)	The total capacity of the volume in GBs.
Device Block Size	The block size of the volume.
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 MBs read each second directly from disks bypassing the cache.
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 61** Snap 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	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 to the disk from the disk director each second.
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.
% 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.

## Spare Disk metrics

The following table lists all metrics that are available for spare disks.

**Table 62** 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.



**Table 62** Spare disk metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
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

**Table 62** Spare disk metrics (continued)

Metric	Description
	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 63** SRP metrics

Metric	Description
Host IOs/sec	The number of host IO operations performed each second, including writes and random and sequential reads.
Host Reads/sec	Host read operations performed each second by the SRP.
Host Writes/sec	Host write operations performed each second by the SRP.
Host MBs/sec	The number of host MBs read each second.
Host MBs Read/sec	The cumulative number of host MBs read per second by the SRP.
Host MBs Written/sec	The cumulative number of host MBs written per second by the SRP.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.

**Table 63** SRP metrics (continued)

<b>Metric</b>	<b>Description</b>
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Req/sec	The number of write requests each second performed by the disk directors to cache.
BE MBs Transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
Response Time (ms)	The average time it takes to satisfy IO requests.
Read RT (ms)	The average time it took the system to serve one read IO for this SRP.
Write RT (ms)	The average time it took the system to serve one write IO for this SRP.
Overall Efficiency Ratio	Ratio of the sum of all TDEVs + snapshot sizes (calculated based on the 128K track size) and the Physical Used Storage (calculated based on the compressed pool track size).
% Snapshot Saved	Percentage savings of the sum of all TDEV Snapshot sizes (at the time of Snapshot creation) and the TDEV Snapshot Allocated Storage.
% Virtual Provisioning Saved	Percentage savings of the TDEV configured storage presented to the hosts and the TDEV Allocated Storage.
Virtual Provisioning Compression Ratio	<p>Ratio of the TDEV Logical Backend Storage (calculated based on the 128K track size) and the TDEV Physical Used Storage (calculated based on the compressed pool track size).</p> <hr/> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p> <hr/>
Virtual Provisioning Efficiency Ratio	Ratio of the TDEV configured storage and the TDEV Logical Backend Storage (calculated based on TDEV allocated tracks without shared unowned).

**Table 63** SRP metrics (continued)

Metric	Description
Virtual Provisioning Shared Ratio	Ratio of the TDEV allocated storage and the TDEV Logical Backend Storage. Backend Storage is the sum of TDEV allocated tracks without counting the TDEV shared unowned tracks.
Snapshot Compression Ratio	<p>Ratio of the RDP Logical Backend Storage (calculated based on the 128K track size) and the RDP Physical Used Storage of the RDP space (calculated based on the compressed pool track size).</p> <hr/> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p> <hr/>
Snapshot Efficiency Ratio	Ratio of the sum of all snapshot sizes and the RDP Logical Backend Storage.
Snapshot Shared Ratio	Ratio of the Snapshot Allocated Storage and the RDP Logical Backend Storage. Sum of all modified tracks on all snapshots / Sum of the RDP size for all snapshot.
Overall Compression Ratio	<p>Ratio of the sum of all TDEVs + snapshot sizes (calculated based on the 128K track size) and the Physical Used Storage (calculated based on the compressed pool track size).</p> <hr/> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p> <hr/>
Total SRP Capacity GB	The total amount of capacity available in all SRPs in GB.
Used SRP Capacity GB	The amount of used capacity in all SRPs in GB.

## Storage Group metrics

The following table lists all metrics that are available for storage groups.

**Table 64** Storage group metrics

<b>Metric</b>	<b>Description</b>
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	The number of IOs the FE port is performing on behalf of the specific storage group.
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	The number of MBs the FE port is processing on behalf of the specific storage group.
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	The number of IOs the FE port is processing on behalf of the specific storage group.
Host Write Misses/sec	The number of host write operations performed each second by the group that were not satisfied from cache.
Host MBs/sec	The cumulative number of host MBs read/writes per second by the group.
Host MBs Read/sec	The cumulative number of host MBs read per second by the group.
Host MBs Written/sec	The cumulative number of host MBs written per second by the group.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read Reqs/sec	The number of read requests each second performed by the disk directors to cache.

**Table 64** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
BE Write Reqs/sec	The number of write requests each second performed by the disk directors to cache.
Read RT (ms)	The calculated average read response time.
Host IO Limit % Time Exceeded	
Write RT (ms)	The calculated average write response time
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.
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.

**Table 64** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
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 sequential.
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.

**Table 64** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the group.
BE MBs Transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	Total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes to make a request by the disk directors to the cache.
BE Disk Read 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	



**Table 64** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	
Write RT Over 32	
Read RT Count 0 to 1	
Response Time (ms)	The calculated average response time.
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.

**Table 64** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	The total allocated storage group capacity.
Total Capacity (GB)	The total capacity of the storage group (GBs).
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.
Compressed Tracks	The number of compressed tracks for the storage group (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the storage group tracks that are compressed tracks (applies only to storage systems running Enginuity OS 5876).
Write paced delay	
Avg. Write Paced Delay (ms)	
Sync Write Reqs/sec	The number of sync write requests managed by this storage group per second.
Sync MBs Sent/sec	The number of sync MBs sent per second via this storage group.
Async Write Reqs/sec	The number of async write requests managed by this storage group per second.
Async MBs Sent/sec	The number of async MBs sent per second via this storage group.
Information Alert Count	Number of information alerts.
Warning Alert Count	Number of warning alerts.
Critical Alert Count	Number of critical alerts.
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	

**Table 64** Storage group metrics (continued)

Metric	Description
RDF Read Hits/sec	
RDF Rewrites/sec	
% Virtual Provisioning Space Saved	
Compression Ratio	<p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p>
QOS Delayed IO Count	The amount of IO delayed that was introduced by enabling QoS services.
QOS Delayed IO Time ( $\mu$ s )	The amount of time IO was delayed that was introduced by enabling QoS services.
QOS Host Delayed IO Count	The amount of Host IO delayed that was introduced by enabling QoS services.
QOS Host Delayed IO Time ( $\mu$ s )	The amount of time IO was delayed that was introduced by enabling QoS services.
Response Time (ms)	The calculated average response time.
PowerPath Average Response Time (ms)	The average response time in milliseconds using PowerPath metrics.
PowerPath Observed Delta RT (ms)	The difference between PowerPath average response time and device/storage group response time.
PowerPath Observed Relative RT (%)	The relative response times between the PowerPath average response time and the device/storage group response time.
Tracks Deduped Not Written	Indicates number of LRU writes deduped and not producing backend IO. This metric will be used for rate calculation.
Writes Broke Dedup Relationships	Indicates number of writes forcing dedup-split (2*128K IO's and allocation expense). This metric will be used for rate calculation.
Duplicate Track Count	Indicates capacity cost of disabling de-dup for a given device. This is used to indicate the number of de-dup track at the instance of the query and not used for rate calculation.

## Storage Group by Pool metrics

The following table lists all metrics that are available for storage groups (by pool).

**Table 65** 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.
Allocated Tracks	
Compressed Tracks	The number of compressed tracks in the storage group (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the total tracks that are compressed (applies only to storage systems running Enginuity OS 5876).
Local Replication Tracks	

## Storage Group by Tier metrics

The following table lists all metrics that are available for storage groups (by tier).

**Table 66** 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 Tracks	The number of allocated tracks in the storage group (by tier).
IO Density	The number of BE requests per GB of disk. (BE Reads + BE Writes) / allocated capacity With FAST moving active extents to higher

**Table 66** Storage group (by tier) metrics (continued)

Metric	Description
	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 (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the storage group tracks that are compressed tracks (applies only to storage systems running Enginuity OS 5876).
Local Replication Tracks	

## Thin Pool metrics

The following table lists all metrics that are available for thin pools.

**Table 67** 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.

**Table 67** Thin pool metrics (continued)

<b>Metric</b>	<b>Description</b>
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	The number of tracks currently in write pending mode for the thin pool.
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)
% 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.
Ingress Tracks	The number of tracks entering the pool.
Egress Tracks	The number of tracks leaving the pool.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
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

**Table 67** Thin pool metrics (continued)

Metric	Description
	success (the IO density on Flash tiers should be higher than the density on SATA tiers.)
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 (KB)	Internal metric.
BE Blocks Written For Copy (KB)	
BE Blocks Read For VLUN Migration	Internal metric.
BE Reads For VLUN Migration	Internal metric.
BE Writes For VLUN Migration	Internal metric.
BE Blocks Written For VLUN Migration	
BE Writes For Rebuild	Internal metric.
BE Prefetched Tracks/sec	
BE Writes For Rebuild (KB)	Internal metric.
BE RDF Copy	Internal metric.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE RDF Copy (MB)	Internal metric.
Total Capacity (GB)	The total capacity of the thin pool (GBs).
Device Block Size	The block size of the volume.
Total Tracks	The number of allocated tracks in the thin pool.
Compressed Tracks	The number of compressed tracks in the thin pool (applies only to storage systems running Enginuity OS 5876) .
Compression Scan Reads	
% Compressed Tracks	The percent of the total tracks that are compressed (applies only to storage systems running Enginuity OS 5876).

## Thin Volume metrics

The following table lists all metrics that are available for thin volumes.

**Table 68** Thin volume metrics

<b>Metric</b>	<b>Description</b>
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	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 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 calculated average read response time.
Write RT (ms)	The calculated average write response time.
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.



**Table 68** Thin volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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 each second by the volume that were misses.

**Table 68** Thin volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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: $100 * (\text{total seq ios per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.

**Table 68** Thin volume metrics (continued)

<b>Metric</b>	<b>Description</b>
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the volume.
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.
% 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 cache to disks.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
Number of RDF Invalid Tracks	
RDF Response Time (ms)	

**Table 68** Thin volume metrics (continued)

<b>Metric</b>	<b>Description</b>
BE RDF Copy	
BE RDF Copy (MB)	
Allocated Capacity (GB)	The allocated capacity of the thin volume (GB).
Total Capacity (GB)	The total capacity of the thin volume (GB).
Device Block Size	The block size of the volume.
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 (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	
Random Writes to WP slots per second	Number of writes to affect the back end.
Random Writes to WB Slots in KB per second.	The amount of writes to affect the back end (KB).
QOS Delayed IO Count	The amount of IO delayed that was introduced by enabling QoS services.
QOS Delayed IO Time ( $\mu$ s )	The amount of time IO was delayed that was introduced by enabling QoS services.
QOS Host Delayed IO Count	The amount of Host IO delayed that was introduced by enabling QoS services.
QOS Host Delayed IO Time ( $\mu$ s )	The amount of time IO was delayed that was introduced by enabling QoS services.
Response time (ms)	The calculated average response time.
Response Time (ms)	The calculated average response time.
PowerPath Average Response Time (ms)	The average response time in milliseconds using PowerPath metrics.

**Table 68** Thin volume metrics (continued)

Metric	Description
PowerPath Observed Delta RT (ms)	The difference between PowerPath average response time and device/storage group response time.
PowerPath Observed Relative RT (%)	The relative response times between the PowerPath average response time and the device/storage group response time.
Tracks Deduped Not Written	Indicates number of LRU writes deduped and not producing backend IO. This metric will be used for rate calculation.
Writes Broke Dedup Relationships	Indicates number of writes forcing dedup-split (2*128K IO's and allocation expense). This metric will be used for rate calculation.
Duplicate Track Count	Indicates capacity cost of disabling de-dup for a given device. This is used to indicate the number of de-dup track at the instance of the query and not used for rate calculation.

## Tier by Storage Group metrics

The following table lists all metrics that are available for tiers (by storage group).

**Table 69** 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	The number of writes per second that were destaged to disk.
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	The size (MBs) of writes per second that were destaged to disk.

**Table 69** Tier (by storage group) metrics (continued)

Metric	Description
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 Tracks	The number of allocated tracks in the storage group.
Compressed Tracks	The number of compressed tracks for the storage group (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the storage group tracks that are compressed tracks (applies only to storage systems running Enginuity OS 5876).
Local Replication Tracks	

## Virtual Pool Tier metrics

The following table lists all metrics that are available for virtual pool tiers.

**Table 70** Virtual pool tier metrics

Metric	Description
Allocated Capacity (GB)	Allocated capacity in GBs.
Allocated Pool Capacity (GB)	The allocated pool capacity in GBs.
Total Tracks	The total number of tracks for the virtual pool.
BE Disk Read RT (ms)	The average time it takes read requests from the disk directors to cache.
BE Read Request 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.
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)
BE MBs Read/sec	The number of MBs read each second by the disk directors.
BE MBs Transferred/sec	The MBs transferred per second to the disk directors.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.

**Table 70** Virtual pool tier metrics (continued)

<b>Metric</b>	<b>Description</b>
BE % Reads	The percent of the back-end IO that were read requests.
BE % Writes	The percent of the back-end IO that were write requests.
BE 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 Req/sec	The number of read requests each second performed by the disk directors to cache.
BE RT (ms)	
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 Write Req/sec	The number of write requests each second performed by the disk directors to cache.
Compressed Tracks	The number of compressed tracks in the thin pool.(applies only to storage systems running Enginuity OS 5876).
BE RDF Copy (MB)	Internal metric.
BE Reads For Copy (KB)	Internal metric.
BE Blocks Read For VLUN Migration	
BE Blocks Written For Copy (KB)	
BE Writes For Rebuild (KB)	Internal metric.
BE Blocks Written For VLUN Migration	
BE Optimize Writes (KB)	Internal metric.
BE Partial Sector Writes (KB)	The number of partial sector writes by the disk directors.
BE RDF Copy	Internal metric.
BE Reads For VLUN Migration	Internal metric.
BE Reads For Copy	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Writes For Rebuild	Internal metric.
BE Writes For VLUN Migration	Internal metric.
BE XOR Reads (KB)	Internal metric.
BE XOR Reads	Internal metric.

**Table 70** Virtual pool tier metrics (continued)

<b>Metric</b>	<b>Description</b>
Device Block Size	The block size of the volume.
Total Capacity (GB)	The total capacity of the virtual pool tier.
Egress Tracks	The number of tracks leaving the pool.
Enabled Pool Capacity (GB)	The enabled pool capacity in GBs.
Ingress Tracks	The number of tracks entering the pool.
IO Density	The number of BE requests per GB of disk.
% Compressed Tracks	The percent of the total tracks that are compressed.
% Reads	The percentage of IO operations that were reads.
% Used Capacity	The percent of the virtual pool tier that is used.
% Write	The percentage of IO operations that were writes.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache per second.
Response Time (ms)	The average time it takes to satisfy IO requests.
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.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
Total Pool Capacity (GB)	The total capacity of the virtual pool.
Used Pool Capacity (GB)	The used pool capacity in GBs.
WP Count	The number of tracks currently in write pending mode.

## Volume metrics

The following table lists all metrics that are available for volumes.

**Table 71** Volume metrics

<b>Metric</b>	<b>Description</b>
Host IOs/sec	The total number of host read IO and write IO operations performed each second by the volume.



**Table 71** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
Host Reads/sec	The total number of host read IO operations performed each second by the Symmetrix volume.
Host Writes/sec	The total number of host write IO operations performed each second by the Symmetrix volume.
Host Hits/sec	The total number of host read IO and write IO operations performed each second by the Symmetrix volume that were immediately satisfied by cache.
Host Read Hits/sec	The total number of host read IO operations performed each second by the Symmetrix volume that were immediately satisfied by cache.
Host Write Hits/sec	The total number of host write IO operations performed each second by the Symmetrix volume that were immediately satisfied by cache.
Host Misses/sec	The total number of host read IO and write IO operations missed each second by the Symmetrix volume.
Host Read Misses/sec	The total number of host read IO operations missed each second by the Symmetrix volume.
Host Write Misses/sec	The total number of host write IO operations missed each second by the Symmetrix volume.
Host MBs/sec	A cumulative number of host MBs read/writes per second.
Host MBs Read/sec	A cumulative number of host MBs read per second.
Host MBs Written/sec	A cumulative number of host MBs written per second.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to the cache.
BE Read Reqs/sec	The number of read requests each second performed by the disk directors to the cache.
Parity Generation Kbytes Read	
Parity Generation Reads	
BE Write Reqs/sec	The number of write requests each second performed by the disk directors to the cache.

**Table 71** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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 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.

**Table 71** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.

**Table 71** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the Symmetrix volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes 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 - 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.

**Table 71** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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	The number of BE requests per GB of disk.
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)	Internal metric.
Kbytes Written for IVTOC	
BE Writes For Rebuild	Internal metric.

**Table 71** Volume metrics (continued)


Metric	Description
BE Writes For Rebuild (KB)	Internal metric.
BE RDF Copy	Internal metric.
BE RDF Copy (MB)	Internal metric.
Allocated Capacity (GB)	The allocated volume capacity.
Total Capacity (GB)	The total capacity of the volume.
Device Block Size	The block size of the volume.
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 (ms)	
Avg. Write Paced Delay (ms)	
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Rewrites/sec	

## Viewing and managing metrics

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

To view and manage metrics:

### Procedure

1. Select  to open the **Settings** panel.
2. In the **Settings** dialog box, select **Performance > Metrics**.
3. 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](#) on page 832.

The following properties display:

#### Metrics

The ID of the performance metric.

#### KPI

An indicator that shows if the metric is a KPI.


4. Select the checkbox for a metric to add it as a KPI or clear the metric's KPI checkbox to no longer have the metric as a KPI.
5. Optional : Click **RESTORE CATEGORY DEFAULTS** to restore the defaults.
6. Click **APPLY**.

## Editing metrics

You can edit metrics to set, or unset, them as KPIs. When a metric is not set as a KPI, it can be viewed under **All** in the Charts view.

To edit metrics:

### Procedure

1. Select .
2. Select **Performance > Settings > Metrics**.
3. Select a metric category from the **Category** menu.
4. Do one of the following:
  - Edit the KPI status of one or more metrics:
    - a. Select or clear the **KPI** checkbox.
    - b. Click **Apply**.
  - Click **Restore Defaults** to restore default metric settings for all metrics in the selected category.

## Metrics

See [Understanding Performance Management](#) on page 518 for an overview of Unisphere Performance Management.

The following table lists all metrics that are available for arrays.

**Table 72** 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 of 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	

**Table 72** Array metrics (continued)

<b>Metric</b>	<b>Description</b>
RDF/A WP Count	The number of RDF/A writes pending.
Number of R2 Active Sessions	
Total Number of RDF Sessions	
Host Reads/sec	The number of host read operations performed each second by all 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 IOs/sec	The total IO from all BE directors to the disks per second.



**Table 72** Array metrics (continued)

<b>Metric</b>	<b>Description</b>
BE Reqs/sec	A data transfer of a read or write between the cache and the director.
BE Read Reqs/sec	A data transfer of a read between the cache and the director.
BE Write Reqs/sec	A data transfer of a write between the cache and the director.
System WP Events/sec	The number of times each second that write activity was heavy enough to use up the system limit set for write tracks occupying cache. When the limit is reached, writes are deferred until data in cache is written to disk.
Device WP Events/sec	The number of times each second that the write-pending limit for a specific 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.
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 MBs read each second directly from disks bypassing the cache.
Avg Optimized Read Miss Size (KB)	

**Table 72** Array metrics (continued)

Metric	Description
FE Balance	Balance information for FE. 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.
Cache Balance	Balance information for cache. 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.
SATA Balance	Balance information for SATA. 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.
FC Balance	Balance information for FC. 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.
EFD Balance	Balance information for EFD. 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.
DX Balance	Balance information for DX. 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.
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: $Sx = 1.9 [(Avg \text{ Queue Depth Range } X - 1) / Avg \text{ Queue Depth Range } X] * (Queue \text{ Depth Count Range } X / Queue \text{ Depth Count Total}) * 100$
DA Balance	Balance information for DA. 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.
Allocated Capacity	The total allocated array capacity.
Critical Alert Count	Number of critical alerts.
Warning Alert Count	Number of warning alerts.
Information Alert Count	Number of information alerts.

**Table 72** Array metrics (continued)

<b>Metric</b>	<b>Description</b>
Disk Utilization	Percentage of time the disk is busy.
RDF Utilization	Percentage of time RDF is busy.
BE Utilization	Percentage of time the BE is busy.
FE Utilization	Percentage of time the FE is busy.
Copy Slot Count	
Overall Efficiency Ratio	Ratio of the sum of all TDEVs + snapshot sizes (calculated based on the 128K track size) and the Physical Used Storage (calculated based on the compressed pool track size).
RDF Balance	Balance information for RDF. 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.
% Snapshot Saved	Percentage savings of the sum of all TDEV Snapshot sizes (at the time of Snapshot creation) and the TDEV Snapshot Allocated Storage.
% Virtual Provisioning Saved	Percentage savings of the TDEV configured storage presented to the hosts and the TDEV Allocated Storage.
Snapshot Compression Ratio	<p>Ratio of the RDP Logical Backend Storage (calculated based on the 128K track size) and the RDP Physical Used Storage of the RDP space (calculated based on the compressed pool track size).</p> <hr/> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p> <hr/>
Snapshot Efficiency Ratio	Ratio of the sum of all snapshot sizes and the RDP Logical Backend Storage.
Snapshot Shared Ratio	Ratio of the Snapshot Allocated Storage and the RDP Logical Backend Storage. Sum of all modified tracks on all snapshots / Sum of the RDP size for all snapshot.
Virtual Provisioning Compression Ratio	Ratio of the TDEV Logical Backend Storage (calculated based on the 128K track size) and the TDEV Physical Used Storage (calculated based on the compressed pool track size).

**Table 72** Array metrics (continued)

Metric	Description
	<p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p>
Virtual Provisioning Efficiency Ratio	Ratio of the TDEV configured storage and the TDEV Logical Backend Storage (calculated based on TDEV allocated tracks without shared unowned).
Virtual Provisioning Shared Ratio	Ratio of the TDEV allocated storage and the TDEV Logical Backend Storage. Backend Storage is the sum of TDEV allocated tracks without counting the TDEV shared unowned tracks.
Overall Compression Ratio	<p>Ratio of the sum of all TDEVs + snapshot sizes (calculated based on the 128K track size) and the Physical Used Storage (calculated based on the compressed pool track size).</p> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p>
% Subscribed Capacity	The subscribed capacity as a percentage of the usable capacity.
% Effective Used Capacity	The percentage of usable capacity that would be used if all compressible data was compressed.
Usable Capacity GB	The total capacity of all Storage Resource Pools (SRP) in GB.
% Metadata Replication Used	The percentage of metadata used for replication.
% Metadata System Used	The percentage of metadata used for the system.
Glacial Track Count	The number of tracks currently in glacially compressed state.
% Meta Data Backend Used	BE TID used percent.
% Meta Data Frontend Used	FE TID used percent.
Snapshot Capacity (GB) + (TB)	Total snapshot capacity for system data.
Snapshot Modified Capacity (GB) + (TB)	Total snapshot modified capacity for system data.
Used Usable Capacity (GB) + (TB)	Total user used capacity.
Capacity Health Score	Capacity Health Score.

**Table 72** Array metrics (continued)

Metric	Description
Hardware Config Health Score	Hardware Config Health Score.

## BE Director (DA) metrics

The following table lists all metrics that are available for back-end directors.

**Table 73** 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.

**Table 73** BE director (DA) metrics (continued)

<b>Metric</b>	<b>Description</b>
MBs Read/sec	The number of reads per second in MBs.
MBs Written/sec	The number of writes per second in MBs.
Clone Copy Read	
Clone Copy Write	
PHCO Rebuild Copy	Internal metric.
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.

**Table 73** BE director (DA) metrics (continued)

Metric	Description
% 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 74** BE emulation metrics

Metric	Description
Total Work Time	
% Busy	The percent of time the director is busy.

## BE Port metrics

The following table lists all metrics that are available for back-end ports.

**Table 75** BE port metrics

Metric	Description
Reads/sec	The number of read operations performed each second by the port.
Writes/sec	The number of write operations performed each second by the port.
Host IOs/sec	The number of host operations performed each second by the port.
MBs Read/sec	The read throughput (MBs) of the port per second.
MBs Written/sec	The write throughput (MBs) of the port per second.
Host MBs/sec	The size of the data transfer from the host in MBs per second.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Speed	Total port capacity.

**Table 75** BE port metrics (continued)

Metric	Description
	Note that in the REST API, for backward compatibility reasons, this metric is named <b>Max Speed Gb/sec.</b>
% Busy	The percent of time that the port is 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 maximum bandwidth.

Utilization, as well as inbound and outbound speed, are measured across all processors, regardless of type.

**Table 76** Board metrics

Metric	Description
Outbound MBs/sec	Outbound speed.
Inbound MBs/sec	Inbound speed.
Max Speed MBs/sec	Maximum speed.
Utilization	Calculated value: $((\text{transfer-rate-in} + \text{transfer-rate-out}) / \text{max-transfer-rate})$

## Cache Partition metrics

The following table lists all metrics that are available for cache partitions.

**Table 77** Cache partition metrics

Metric	Description
Acquired Slot Count	
Age Non WP Slots	
Avg Age Given Destage	
Avg Age of Write to Non WP Slot	$(\text{age non wp slots}) / ((\text{writes to all non wp slots}) - (\text{writes to young non wp slots}))$
Cache Age GT 10 Min	
Cache Age GT 1 hour	
Cache Age GT 1 Min	
Cache Hit Ages 1	
Cache Hit Ages 2	
Cache Hit Ages 3	
Cache Hit Ages 4	



**Table 77** Cache partition metrics (continued)

<b>Metric</b>	<b>Description</b>
Cache Hit Ages 5	
Cache Hit Ages 6	
Cache Hit Ages 7	
Cache Hit Ages 8	
Cache Slot Ages 1	
Cache Slot Ages 2	
Cache Slot Ages 3	
Cache Slot Ages 4	
Cache Slot Ages 5	
Cache Slot Ages 6	
Cache Slot Ages 7	
Cache Slot Ages 8	
Cache Slots Used	
Destage Slot Age	
Destaged Slot Count	
Donation Give Count	
Donation Take Count	
Donation Time	
Host IOs/sec	Host operations performed each second by the cache partition.
Local WP Count	The number of write pending slots waiting to be de-staged to disk on the local system. The value should be less than or equal to the system write pending count.
% Max Cache	Maximum slot allocation for a partition.
Host MBs/sec	The number of host MBs written and read each second.
% Min Cache	Minimum slot allocation for a partition
% 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 percent of IO operations that were immediately satisfied from cache.
Remote WP count	The number of write pending slots waiting to be de-staged to disk on the remote system.

**Table 77** Cache partition metrics (continued)

Metric	Description
	The value should be less than or equal to the system write pending count.
% Cache Used	The percent of the cache partition that is used.
% Target Cache	Percent of total cache allocated to this partition.
Total Replace Slots	
WP Count	The number of cache partition slots that are write pending.
WP Limit	The maximum number of write-pending slots.
Writes to All Non WP Slots	
Writes to Young Non WP Slots	

## DATA Volume metrics

The following table lists all metrics that are available for DATA volumes.

**Table 78** DATA volume metrics

Metric	Description
Allocated Capacity (GB)	The total allocated volume capacity.
BE MBs Read/sec	The number of MBs read by the disk directors from the volume each second.
BE MBs Transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)
BE MBs Written/sec	The number of MBs written to the volume from the disk director each second.
BE % Reads	The percent of the back-end IO that were read requests.
BE % Writes	The percent of the back-end IO that were write requests.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the 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 the cache.
Kbytes Written for IVTOC	
BE Reads For Copy (KB)	Internal metric.
Kbytes Read for Rebuild	

**Table 78** DATA volume metrics (continued)

<b>Metric</b>	<b>Description</b>
BE Reads For VLUN Migration (KB)	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Writes For Rebuild (KB)	Internal metric.
BE Writes For VLUN Migration (KB)	Internal metric.
BE Optimize Writes (KB)	Internal metric.
BE Partial Sector Writes (KB)	Internal metric.
Reads for Rebuild	
BE Reads For VLUN Migration	Internal metric.
BE Reads For Copy	Internal metric.
BE Writes For Copy	Internal metric.
BE Writes For Rebuild	Internal metric.
BE Writes For VLUN Migration	Internal metric.
BE XOR Reads (KB)	Internal metric.
BE XOR Reads	Internal metric.
Device Block Size	The block size of the volume.
Total Capacity (GB)	The total capacity of the volume (GBs).
IO Density	The number of BE requests per GB of disk.
Parity Generation Kbytes Read	
Parity Generation Reads	
BE Prefetched MBs/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.
BE Reqs/sec	The number of read/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.
Writes for IVTOC	

## Database metrics

The following table lists all metrics that are available for databases.

**Table 79** Database metrics

<b>Metric</b>	<b>Description</b>
Host IOs/sec	Host operations performed each second by the group.

**Table 79** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
Writes for IVTOC	
RDF Writes/sec	Total RDF writes per second for the database.
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.
Kbytes written for IVTOC	
Read RT (ms)	The average time that it took the system to serve one read command.

**Table 79** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
Write RT (ms)	The average time that it took the system to serve one write command.
Read Miss RT (ms)	The average time a read miss operation was performed by the group. 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 group. 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	Percentage of IO operations that were reads.
% Writes	Percentage of IO operations that were writes.
% Read Hit	The percentage of read operations performed that were immediately satisfied by cache.
% Write Hit	The percentage of write operations performed that were immediately satisfied by cache.
% Read Miss	The percentage of read miss operations performed, 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 percentage of write miss operations performed, 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.
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 that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed each that were misses.

**Table 79** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
Seq Write Hits/sec	The number of sequential write operations performed by the volume that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by 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: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.

**Table 79** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes read requests from the disk directors to cache.
BE Disk Read 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.
Parity Generation Kbytes Read	
Parity Generation Reads	
% Hit	The percent of operations that were immediately satisfied by cache.
% Miss	The percent of operations that were misses. A miss occurs when the operation cannot be immediately satisfied from cache because the data is not there or the operation has to wait while data is destaged from cache to disks.

**Table 79** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
% Random Reads	The percent of read commands from a host not identified as part of a sequential stream.
% Random Writes	The percent of write commands from a host not identified as part of a sequential stream.
BE % Reads	The percent of read operations from the back-end directors that were immediately satisfied by cache.
BE % Writes	The percent of write operations from the back-end directors that were immediately satisfied by cache.
Write RT Count 5 to 8	
Write RT Count 8 to 14	
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)	The average response time for the reads and writes.
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	Number of information alerts.
Critical Alert Count	Number of critical alerts.
Warning Alert Count	Number of information alerts.
Allocated Capacity (GB)	The total allocated database capacity.



**Table 79** Database metrics (continued)

<b>Metric</b>	<b>Description</b>
BE Partial Sector Writes (KB)	The KBs of back-end IO that were partial sector writes.
IO Density	The number of BE requests per GB of disk.
BE 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 Reads For Rebuild (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.
Reads For Rebuild	
Total Capacity (GB)	The total capacity of the database.
Device Block Size	The block size of the volume.
BE Prefetched MBs/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.
RDF Read Hits/sec	
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.
RDF MBs Read/sec	
RDF Rewrites/sec	
RDF MBs Written/sec	
Avg Optimized Read Miss Size (KB)	

**Table 79** Database metrics (continued)

Metric	Description
Write Paced Delay (ms)	

## Database by Pool metrics

The following table lists all metrics that are available for databases by pool.

**Table 80** Database 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 allocated pool 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	The number of writes per second that were destaged to disk.
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.
Destage Write MBs/sec	The size (MBs) of writes per second that were destaged to disk.
Prefetched Tracks/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.
Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
Total Tracks	The number of allocated tracks in the database by pool.
Compressed Tracks	The number of compressed tracks in the database by pool (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the total tracks that are compressed (applies only to storage systems running Enginuity OS 5876).

## Device Group metrics

The following table lists all metrics that are available for device groups.

**Table 81** 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	Total RDF writes per second for the device group.
Host Writes/sec	The number of host write operations performed each second by the group.
Host Hits/sec	The number of host read/write operations performed each second by the group that were immediately satisfied from cache.
Host Read Hits/sec	The number of host read operations performed each second by the group that were immediately satisfied from cache.
Host Write Hits/sec	The number of host write operations performed each second by the group that were immediately satisfied from cache.
Host Misses/sec	The number of host read/write operations performed each second by the group that could not be satisfied from cache.
Host Read Misses/sec	The number of host read operations performed each second by the group that were not satisfied from cache.
Host Write Misses/sec	The number of host write operations performed each second by the group that were not satisfied from cache.
Host MBs/sec	The cumulative number of host MBs read/writes per second by the group.
Host MBs Read/sec	The cumulative number of host MBs read per second by the group.
Host MBs Written/sec	The cumulative number of host MBs written per second by the group.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read Reqs/sec	The number of read requests each second performed by the disk directors to cache.

**Table 81** Device group metrics (continued)

<b>Metric</b>	<b>Description</b>
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 satisfied from cache.

**Table 81** Device group metrics (continued)

<b>Metric</b>	<b>Description</b>
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	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.

**Table 81** Device group metrics (continued)

<b>Metric</b>	<b>Description</b>
% Random IO	The percent of IO operations that were random.
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
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 requests.
BE % Writes	The percent of the back-end IO that were write requests.
RDF Response Time (ms)	

**Table 81** Device group metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	The total capacity of the device group.
Allocated Capacity (GB)	The allocated capacity of the device group.
Device Block Size	The block size of the volume.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
RDF Rewrites/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Read Hits/sec	
Write Paced Delay (ms)	
Avg. Write Paced Delay (ms)	

## Disk metrics

The following table lists all metrics that are available for disks.

**Table 82** Disk metrics

<b>Metric</b>	<b>Description</b>
% 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.
Avg Hypers per Seek	The average number of hypervolumes that the disk head crossed during one seek.
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 83** 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 84** Disk group metrics

Metric	Description
Avg Read Size (KB)	Calculated value: (Kbytes read per sec/total reads per sec)

**Table 84** Disk group metrics (continued)

Metric	Description
Avg Write Size (KB)	Calculated value: (Kbytes written per sec/total writes per sec)
IO Density	The number of BE requests per GB of disk.
IOs/sec	The number of IO commands to the disk.
MBs/sec	The total IO (reads and writes) per second in MBs.
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.
% Busy	The percent of time that the disk is busy serving IOs.
% Idle	The percent of time the disk is idle.
% Free Capacity	The percent of the disk that is free.
% Used Capacity	The percent of the disk group capacity that is allocated.
Reads/sec	The number of host reads per second for the disk.
Writes/sec	The number of host writes per second for the disk.
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.
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 each second.
Total Capacity (GB)	The total capacity of the disk (GBs).
Used Capacity (GB)	The total used capacity of the disk (GBs).

## Disk Group tier metrics

The following table lists all metrics that are available for disk group tiers.

**Table 85** 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.

**Table 85** Disk group tier metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
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 + 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 86** Disk technology metrics

<b>Metric</b>	<b>Description</b>
% 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.

**Table 86** Disk technology metrics (continued)

Metric	Description
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.)

## DSE Pool metrics

The following table lists all metrics that are available for DSE pools.

**Table 87** 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 88** DX emulation metrics

Metric	Description
Total Work Time	
% Busy	The percent of time the directory is busy.

## DX Port metrics

The following table lists all metrics that are available for DX ports.

**Table 89** DX 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	
% Busy	

## EDS Director metrics

The following table lists all metrics that are available for EDS directors.

**Table 90** EDS director metrics

Metric	Description
Random Read Misses	The number of random read IOs that were misses.
Random Read Misses Mbytes	The number of random read IOs that were misses in MBs.
Random Write Misses	The number of random write IOs that were misses.
Random Write Misses Mbytes	The number of random write IOs that were misses in MBs.
% Busy	The percent of time the director is busy.
Glacial Tracks In	Ingress to glacial tier.

**Table 90** EDS director metrics (continued)

Metric	Description
Glacial Tracks Out	Egress from glacial tier.

## EDS Emulation metrics

The following table lists all metrics that are available for EDS emulations.

**Table 91** EDS director metrics

Metric	Description
Total Work Time	
% Busy	The percent of time the directory is busy.

## External Director metrics

The following table lists all metrics that are available for external directors.

**Table 92** External director 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.

**Table 92** External director metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
Clone Copy Read	
Clone Copy Write	
PHCO Rebuild Copy	Internal metric.
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.



**Table 92** External director metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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.

## External Disk metrics

The following table lists all metrics that are available for external disks.

**Table 93** External disk metrics

<b>Metric</b>	<b>Description</b>
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.

**Table 93** External disk metrics (continued)

<b>Metric</b>	<b>Description</b>
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 94** External disk group metrics

<b>Metric</b>	<b>Description</b>
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)	The average response time for the reads and writes.
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.
Members	The number of members in the external disk group.

## FAST VP Policy metrics

The following table lists all metrics that are available for FAST VP policies.

**Table 95** 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 96** FE director metrics

Metric	Description
% Busy	The percent of time the director 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

**Table 96** FE director metrics (continued)

<b>Metric</b>	<b>Description</b>
	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.
% 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.

**Table 96** FE director metrics (continued)

<b>Metric</b>	<b>Description</b>
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	The percent of time that core 0 is busy serving IOs.
Syscall Time/sec	Total time spent processing all system calls.
Avg RDF Write Time (ms)	
Avg WP Disconnect Time (ms)	
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	The percent of time that core 1 is busy serving IOs.
% 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.

**Table 96** FE director metrics (continued)

Metric	Description
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 1toRead RT Over 64	These buckets show the distribution of the total average read response time.
Write RT 0 to 1toWrite RT Over 64	These buckets show the distribution of the total average write response time.
Host IO Limit IOs/sec	The number of IOs the FE director is processing.
Host IO Limit MBs/sec	The number of MBs the FE port is processing.
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: $S_x = 1..9 [(Avg Queue Depth Range X - 1) / Avg Queue DepthRange X] * (Queue Depth Count Range X / Queue Depth Count Total) * 100$
CCWS Per IO	
Avg Optimized Read Miss Size (KB)	
Optimized MBs Read Miss/sec	Number of host MBs read each second directly from disks bypassing the cache.
Optimized Read Misses/sec	Number of read requests each second performed directly from disks bypassing the cache.
ORS ASync IOs/sec	
ORS ASync MBs/sec	
ORS Sync MBs/sec	
ZHPF reads	
ZHPF Reads MBs	
ZHPF Writes	
ZHPF Writes MBs	

**Table 97** FE director metrics

Metric	Description
% Busy	The percent of time the director is busy.

**Table 97** FE director metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
% Read Reqs	The percent of read requests out of host commands for data transfer.
% Write Reqs	The percent of write requests out of host commands for data transfer.



**Table 97** FE director metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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	The percent of time that core 0 is busy serving IOs.
Syscall Time/sec	Total time spent processing all system calls.
Avg RDF Write Time (ms)	
Avg WP Disconnect Time (ms)	

**Table 97** FE director metrics (continued)

<b>Metric</b>	<b>Description</b>
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	The percent of time that core 1 is busy serving IOs.
% 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 1toRead RT Over 64	These buckets show the distribution of the total average read response time.
Write RT 0 to 1toWrite RT Over 64	These buckets show the distribution of the total average write response time.
Host IO Limit IOs/sec	The number of IOs the FE director is processing.
Host IO Limit MBs/sec	The number of MBs the FE port is processing.
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: $S_x = 1..9 [(Avg Queue Depth Range X - 1) / Avg Queue Depth Range X] * (Queue Depth Count Range X / Queue Depth Count Total) * 100$
CCWS Per IO	
Avg Optimized Read Miss Size (KB)	
Optimized MBs Read Miss/sec	Number of host MBs read each second directly from disks bypassing the cache.
Optimized Read Misses/sec	Number of read requests each second performed directly from disks bypassing the cache.

**Table 97** FE director metrics (continued)

Metric	Description
ORS ASync IOs/sec	
ORS ASync MBs/sec	
ORS Sync MBs/sec	
ZHPF reads	
ZHPF Reads MBs	
ZHPF Writes	
ZHPF Writes MBs	
EPL Busy	
SPA Busy	

## FE Emulation metrics

The following table lists all metrics that are available for FE emulations.

**Table 98** FE emulation metrics

Metric	Description
Total Work Time	
% Busy	The percent of time the directory is busy.

## FE Port metrics

The following table lists all metrics that are available for FE ports.

**Table 99** FE port metrics

Metric	Description
Reads/sec	The average number of host reads performed per second.
Writes/sec	The average number of host writes performed per second.
Host IOs/sec	The number of IOs the port is performing on behalf of the specific SG.
MBs Read/sec	The throughput in MBs read per second.
MBs Written/sec	The throughput in MBs written per second.
Host MBs/sec	The number of MBs the port is processing on behalf of the specific SG.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Speed Gb/sec	

**Table 99** FE port metrics (continued)

Metric	Description
Max Speed Gb/sec	
% Busy	The percent of time the port is busy.
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.
Response Time (ms)	The average response time for the reads and writes.
iSCSI Checksum Error Count	The percent of time that the port is busy.

## FE Port - FE metrics

The following table lists all metrics that are available for FE ports (FE).

**Table 100** FE port (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	

## FE Port - SG metrics

The following table lists all metrics that are available for FE ports (SG).

**Table 101** FE port (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	

## FICON Emulation metrics

The following table lists all metrics that are available for FICON emulations.

**Table 102** FICON emulation metrics

Metric	Description
% Busy	The percent of time the FICON channel is busy.

## FICON Emulation Threads metrics

The following table lists all metrics that are available for FICON emulation threads.

**Table 103** FICON emulation threads metrics

Metric	Description
% Busy	The percent of time the FICON emulation threads were busy.
% Idle	The percent of time the FICON emulation threads were idle.

## FICON Port Threads metrics

The following table lists all metrics that are available for FICON emulations.

**Table 104** FICON port threads metrics

Metric	Description
% Busy	The percent of time the FICON port was busy.
% Idle	The percent of time the FICON port was idle.

## Host metrics

The following table lists all metrics that are available for hosts.

**Table 105** Host metrics

Metric	Description
Host IOs/sec	The total number of host read IO and write IO operations performed each second by the host.
Host Reads/sec	The total number of host read IO operations performed each second by the host.
Host Writes/sec	The total number of host write IO operations performed each second by the host.
Host MBs/sec	The size of the data transfer from the host in MBs per second.
Read RT (ms)	The average time it took the system to serve one read IO for this host.

**Table 105** Host metrics (continued)

Metric	Description
Write RT (ms)	The average time it took the system to serve one write IO for this host.
Syscall Count/sec	The total number of calls seen by this host.
Host MBs Read/sec	Cumulative number of host MBs read per second.
Host MBs Written/sec	Cumulative number of host MBs written per second.
Response Time (ms)	The average response time for the reads and writes.

### Host IO Limit by FE metrics

The following table lists all metrics that are available for host IO limits (by FE).

**Table 106** 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 107** 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 108** IM director metrics

<b>Metric</b>	<b>Description</b>
% Busy	The percent of time the director is 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 109** IM emulations metrics

<b>Metric</b>	<b>Description</b>
Total Work Time	
% Busy	The percent of time the directory is busy.

## Initiator metrics

The following table lists all metrics that are available for initiators.

**Table 110** Initiator metrics

<b>Metrics</b>	<b>Description</b>
Host IOs/sec	A host command for data transfer.
Host Reads/sec	The number of host read operations performed each second by the initiator.
Host Writes/sec	The number of host write operations performed each second by the initiator.
Host MBs/sec	A cumulative number of host MBs read/writes per second.
Read RT (ms)	The average time it took the system to serve one read IO for this initiator.
Write RT (ms)	The average time it took the system to serve one write IO for this initiator.
Syscall Count/sec	The total number of calls seen by this initiator.
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.
Response Time (ms)	The average response time for the reads and writes.

## Initiators by Port metrics

The following table lists all metrics that are available for initiators (by port).

**Table 111** Initiators (by port) metrics

Metrics	Description
Host IOs/sec	A host command for data transfer.
Host Reads/sec	The number of host read operations performed each second by the initiator (by port).
Host Writes/sec	The number of host write operations performed each second by the initiator (by port).
Host MBs/sec	A cumulative number of host MBs read/writes per second.
Read RT (ms)	The average time it took the system to serve one read IO for this initiator (by port).
Write RT (ms)	The average time it took the system to serve one write IO for this initiator (by port).
Syscall Count/sec	The total number of calls seen by this initiator (by port).
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.
Response Time (ms)	The average response time for the reads and writes.

## IP Interface metrics

The following table lists all metrics that are available for IP interfaces.

**Table 112** IP interface metrics

Metric	Description
Packets/sec	
MBs/sec	
Duplicate Acks Received/sec	
TCP Retransmits/sec	
Reads/sec	The number of read operations performed each second.
Writes/sec	The number of write operations performed each second.
Host IOs/sec	The number of host operations performed each second.
Total Read Time	Total time spent performing reads.
Total Write Time	Total time spent performing writes.



**Table 112** IP interface metrics (continued)

Metric	Description
Response Time (ms)	The average response time for the reads and writes.

## iSCSI Target metrics

The following table lists all metrics that are available for iSCSI targets.

**Table 113** iSCSI target metrics

Metric	Description
Packets/sec	
MBs/sec	
Duplicate Acks Received/sec	
TCP Retransmits/sec	
Reads/sec	The number of read operations performed each second.
Writes/sec	The number of write operations performed each second.
Host IOs/sec	The number of host operations performed each second.
Total Read Time	Total time spent performing reads.
Total Write Time	Total time spent performing writes.
Response Time (ms)	The average response time for the reads and writes.

## Masking View metrics

The following table lists all metrics that are available for masking views.

**Table 114** Masking view metrics

Metric	Description
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.
Response Time (ms)	The average response time for the reads and writes.
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.

**Table 114** Masking view metrics (continued)

Metric	Description
Capacity (GB)	The total capacity of the masking group in GBs.

## Metas metrics

The following table lists all metrics that are available for metas.

**Table 115** 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.

**Table 115** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
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

**Table 115** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
	volume. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.
WP Count	The number of tracks currently in write pending mode for the volume.
Seq IOs/sec	The number of IO operations performed each second that were sequential.
Seq Reads/sec	The number of read IO operations performed each second that were sequential.
Seq Writes/sec	The number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	The number of sequential read operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed each second by the Symmetrix volume that were misses.
Seq Write Hits/sec	The number of sequential write operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by the Symmetrix volume that were misses.
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.

**Table 115** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the Symmetrix volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.

**Table 115** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
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	The number of BE requests per GB of disk.
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.

**Table 115** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	The allocated capacity of the meta.
Total Capacity (GB)	The total capacity of the meta.
Device Block Size	The block size of the volume.
XtremSW Cache Read Hits/sec	The cumulative number of reads that were hits by XtremSW Cache per second.
XtremSW Cache Write Hits/sec	The cumulative number of writes that were hits by XtremSW Cache per second.
XtremSW Cache Reads/sec	The average time it took XtremSW Cache to serve one read.
XtremSW Cache Writes/sec	The average time it took XtremSW Cache to serve one write.
XtremSW Cache IOs/sec	The average time it took XtremSW Cache to serve one IO.
XtremSW Cache Skipped IOs/sec	The number of IOs that were skipped by XtremSW Cache per second.
XtremSW Cache Dedup Hits/sec	The number of XtremSW Cache Dedup hits per second.
XtremSW Cache Dedup Reads/sec	The number of XtremSW Cache Dedup reads per second.
XtremSW Cache Dedup Writes/sec	The number of XtremSW Cache Dedup writes per second.
% XtremSW Cache Reads	The percent of XtremSW Cache IO that were reads.

**Table 115** Metas metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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	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 (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 116** 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 performed by the disk directors to the cache.
Parity Generation Reads	

**Table 116** Pool-bound volumes metrics (continued)

<b>Metric</b>	<b>Description</b>
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 mode for the volume.
Seq IOs/sec	The number of IO operations performed each second that were sequential.

**Table 116** Pool-bound volumes metrics (continued)

<b>Metric</b>	<b>Description</b>
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)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$

**Table 116** Pool-bound volumes metrics (continued)

Metric	Description
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the Symmetrix volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes read requests from the disk directors to cache.
BE Disk Read 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.

**Table 116** Pool-bound volumes metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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	The number of BE requests per GB of disk.
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	

**Table 116** Pool-bound volumes metrics (continued)

Metric	Description
BE Writes For VLUN Migration (KB)	Internal metric.
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)	The allocated pool capacity in GBs.
Total Capacity (GB)	The total capacity of the pool in GBs.
Device Block Size	The block size of the volume.
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 MBs read each second directly from disks bypassing the cache.
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 117** Pool by storage group metrics

Metric	Description
BE Reqs/sec	The number of data transfers between cache and the pool by storage group.
Allocated Capacity	The allocated capacity for the pool by storage group in GBs.
IO Density	The number of BE requests per GB of disk. (BE Reads + destaged writes) / capacity

**Table 117** Pool by storage group metrics (continued)

Metric	Description
BE Read Reqs/sec	A data transfer of a read between the cache and the director.
Destage Writes/sec	The number of writes per second that were destaged to disk.
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.
Destage Write MBs/sec	The size (MBs) of writes per second that were destaged to disk.
Prefetched Tracks/sec	The number of tracks per second prefetched from disk to cache upon detection of a sequential read stream.
Pretetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
Allocated Tracks	The number of allocated tracks in the pool.
Compressed Tracks	The number of compressed tracks (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the total tracks that are compressed (applies only to storage systems running Enginuity OS 5876).
Local Replication Tracks	

## Port Group metrics

The following table lists all metrics that are available for port groups.

**Table 118** Port group metrics

Metric	Description
Reads/sec	The average number of host reads performed per second.
Writes/sec	The average number of host writes performed per second.
Host IOs/sec	The number of host IO operations performed each second, including writes and random and sequential reads.
MBs Read/sec	The number of reads per second in MBs.
MBs Written/sec	The number of writes per second in MBs.
Host MBs/sec	The number of host MBs read each second.

**Table 118** Port group metrics (continued)

Metric	Description
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
% Busy	The percent of time the port group is busy.

## RDF Director metrics

The following table lists all metrics that are available for RDF directors.

**Table 119** 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.
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	
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.



**Table 119** RDF director metrics (continued)

Metric	Description
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 (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 MBs Sent/sec	The number of async MBs sent per second via this RDF director.
% Compressed MBs Sent and Received/sec	Calculated value: ((Compressed MBs Sent and Received/sec / MBs Sent and Received/sec) * 100)
% Compressed MBs Received/sec	Calculated value: ((Compressed MBs Received/sec / MBs Received/sec) * 100)
% Compressed MBs Sent/sec	Calculated value: ((Compressed MBs Sent/sec / MBs Sent/sec) * 100)
Compressed MBs Sent and Received/sec	The number of compressed MBs sent and received by this RDF director.
Compressed MBs Received/sec	The number of compressed MBs/sec received by this RDF director.
Compressed MBs Sent/sec	The number of compressed MBs/sec sent by this RDF director.
Number of Compressed Links	The number of compressed links used by this RDF director.
Number of Links	The number of links used by this RDF director.
Reqs/sec	The average number of requests performed by the director per second.

## RDF Emulation metrics

The following table lists all metrics that are available for RDF emulations.

**Table 120** RDF emulation metrics

Metric	Description
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 121** RDF port metrics

Metric	Description
Reads/sec	The number of read operations performed by the port per second.
Writes/sec	The number of write operations performed by the port per second.
Host IOs/sec	The number of host operations performed each second by the port.
MBs Read/sec	The read throughput (MBs) of the port per second.
MBs Written/sec	The write throughput (MBs) of the port per second.
Host MBs/sec	The size of the data transfer from the host in MBs per second.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Speed Gb/sec	
Max Speed Gb/sec	
% Busy	The percent of time that the port is busy.

## RDF/A Group metrics

The following table lists all metrics that are available for RDF/A groups.

**Table 122** 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.)

**Table 122** RDF/A group metrics (continued)

<b>Metric</b>	<b>Description</b>
RDF Reads/sec	
Avg Cycle Size	The average size of all the cycles in RDFA sessions.
Avg IO Service Time (ms)	The average time the director takes to serve IO.
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.
System WP Limit	The maximum number of write-pending slots for the entire storage system.
System WP Count	The number of system cache slots that are write pending.
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.
DSE Used Tracks	The active cycle spilled count plus the inactive cycle spilled count.

**Table 122** RDF/A group metrics (continued)

Metric	Description
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).
RDF/A Session Indicator	Indicates if group is RDFA (0/1).
RDF/A Active Indicator	Indicates if RDF/A group is active (0/1).
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.
R1 to R2 Lag Time	The time that R2 is behind R1 (RPO). This is calculated as the last cycle time plus the time since last switch. In a regular case, the cycles switch every ~30 sec and the samples are taken every few minutes, therefore this counter may not show very significant data, however, when cycles elongate beyond the sample time, this counter can help indicate an estimate of the RPO.
Num Devs in Session	The number of 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

**Table 122** RDF/A group metrics (continued)

Metric	Description
	one of the volumes reaches its WP limit. The importance of this counter is that it shows one of the common reasons for an increase in the cycle time.
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.
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 number 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)
RDF Rewrites/sec	
RDF MBs Written/sec	
RDF MBs Read/sec	

**Table 122** RDF/A group metrics (continued)

Metric	Description
RDF Read Hits/sec	
Total Tracks Spilled Back	
Total Tracks Spilled Over	
Total Tracks Spilled Over Used	

## RDF/S Group metrics

The following table lists all metrics that are available for RDF/S groups.

**Table 123** RDF/S group metrics

Metric	Description
Host IOs/sec	The number of host operations performed each second by the RDF/S Group.
Host Reads/sec	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.

**Table 123** RDF/S group metrics (continued)

<b>Metric</b>	<b>Description</b>
Host MBs Read/sec	The cumulative number of host MBs read per second by the group.
Host MBs Written/sec	The cumulative number of host MBs written per second by the group.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read 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.

**Table 123** RDF/S group metrics (continued)

<b>Metric</b>	<b>Description</b>
Seq IOs/sec	Number of IO operations performed each second that were sequential.
Seq Reads/sec	Number of read IO operations performed each second that were sequential.
Seq Writes/sec	Number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	Number of sequential read operations performed each second by the group that were immediately satisfied from cache.
Seq Read Miss/sec	Number of sequential read operations performed each second by the group that were misses.
Seq Write Hits/sec	Number of sequential write operations performed each second by the group that were immediately satisfied from cache.
Seq Write Misses/sec	Number of sequential write operations performed each second by the group that were misses.
Random IOs/sec	IOs from a host not identified as part of a sequential stream.
Random Reads/sec	Read IO commands from a host not identified as part of a sequential stream.
Random 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	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$



**Table 123** RDF/S group metrics (continued)

Metric	Description
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
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.

**Table 123** RDF/S group metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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.
RDF Response Time (ms)	
Response Time (ms)	The average time it takes to satisfy IO requests.
IO Density	The number of partial sector writes by the disk directors.
BE Partial Sector Writes (KB)	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)	The total capacity of the device group.
Device Block Size	The block size of the volume.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
RDF MBs Read/sec	
RDF Read Hits/sec	

**Table 123** RDF/S group metrics (continued)

Metric	Description
RDF MBs Written/sec	
RDF Rewrites/sec	
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 124** 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.

**Table 124** SAVE volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
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

**Table 124** SAVE volume metrics (continued)

Metric	Description
	operation had to wait while data was destaged from cache to the disks.
% Write Miss	The percent of write miss operations performed each second by the Symmetrix volume. A miss occurs when the write operation had to wait while data was destaged from cache to the disks.
WP Count	The number of tracks currently in write pending mode for the volume.
Seq IOs/sec	The number of IO operations performed each second that were sequential.
Seq Reads/sec	The number of read IO operations performed each second that were sequential.
Seq Writes/sec	The number of write IO operations performed each second that were sequential.
Seq Read Hits/sec	The number of sequential read operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Read Miss/sec	The number of sequential read operations performed each second by the Symmetrix volume that were misses.
Seq Write Hits/sec	The number of sequential write operations performed each second by the Symmetrix volume that were immediately satisfied from cache.
Seq Write Misses/sec	The number of sequential write operations performed each second by the Symmetrix volume that were misses.
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.

**Table 124** SAVE volume metrics (continued)

<b>Metric</b>	<b>Description</b>
Random Write Misses/sec	The number of random write IOs that were misses.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the Symmetrix volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.

**Table 124** SAVE volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	
IO Density	The number of BE requests per GB of disk.
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.

**Table 124** SAVE volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	Internal metric.
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)	The allocated volume capacity in GBs.
Total Capacity (GB)	The total capacity of the volume in GBs.
Device Block Size	The block size of the volume.
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 MBs read each second directly from disks bypassing the cache.
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 125** Snap 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	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 to the disk from the disk director each second.
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.
% 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.

## Spare Disk metrics

The following table lists all metrics that are available for spare disks.

**Table 126** 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.

**Table 126** Spare disk metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
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.

**Table 126** Spare disk metrics (continued)

Metric	Description
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 127** SRP metrics

Metric	Description
Host IOs/sec	The number of host IO operations performed each second, including writes and random and sequential reads.
Host Reads/sec	Host read operations performed each second by the SRP.
Host Writes/sec	Host write operations performed each second by the SRP.
Host MBs/sec	The number of host MBs read each second.
Host MBs Read/sec	The cumulative number of host MBs read per second by the SRP.
Host MBs Written/sec	The cumulative number of host MBs written per second by the SRP.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE Write Req/sec	The number of write requests each second performed by the disk directors to cache.
BE MBs Transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)

**Table 127** SRP metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
Response Time (ms)	The average time it takes to satisfy IO requests.
Read RT (ms)	The average time it took the system to serve one read IO for this SRP.
Write RT (ms)	The average time it took the system to serve one write IO for this SRP.
Overall Efficiency Ratio	Ratio of the sum of all TDEVs + snapshot sizes (calculated based on the 128K track size) and the Physical Used Storage (calculated based on the compressed pool track size).
% Snapshot Saved	Percentage savings of the sum of all TDEV Snapshot sizes (at the time of Snapshot creation) and the TDEV Snapshot Allocated Storage.
% Virtual Provisioning Saved	Percentage savings of the TDEV configured storage presented to the hosts and the TDEV Allocated Storage.
Virtual Provisioning Compression Ratio	<p>Ratio of the TDEV Logical Backend Storage (calculated based on the 128K track size) and the TDEV Physical Used Storage (calculated based on the compressed pool track size).</p> <hr/> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p> <hr/>
Virtual Provisioning Efficiency Ratio	Ratio of the TDEV configured storage and the TDEV Logical Backend Storage (calculated based on TDEV allocated tracks without shared unowned).
Virtual Provisioning Shared Ratio	Ratio of the TDEV allocated storage and the TDEV Logical Backend Storage. Backend Storage is the sum of TDEV allocated tracks without counting the TDEV shared unowned tracks.
Snapshot Compression Ratio	Ratio of the RDP Logical Backend Storage (calculated based on the 128K track size) and the RDP Physical Used Storage of the RDP

**Table 127** SRP metrics (continued)

Metric	Description
	<p>space (calculated based on the compressed pool track size).</p> <hr/> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p> <hr/>
Snapshot Efficiency Ratio	Ratio of the sum of all snapshot sizes and the RDP Logical Backend Storage.
Snapshot Shared Ratio	Ratio of the Snapshot Allocated Storage and the RDP Logical Backend Storage. Sum of all modified tracks on all snapshots / Sum of the RDP size for all snapshot.
Overall Compression Ratio	<p>Ratio of the sum of all TDEVs + snapshot sizes (calculated based on the 128K track size) and the Physical Used Storage (calculated based on the compressed pool track size).</p> <hr/> <p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p> <hr/>
Total SRP Capacity GB	The total amount of capacity available in all SRPs in GB.
Used SRP Capacity GB	The amount of used capacity in all SRPs in GB.

## Storage Group metrics

The following table lists all metrics that are available for storage groups.

**Table 128** 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.

**Table 128** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
Host IO Limit Exceeded Secs	The number of IOs the FE port is performing on behalf of the specific storage group.
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	The number of MBs the FE port is processing on behalf of the specific storage group.
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	The number of IOs the FE port is processing on behalf of the specific storage group.
Host Write Misses/sec	The number of host write operations performed each second by the group that were not satisfied from cache.
Host MBs/sec	The cumulative number of host MBs read/writes per second by the group.
Host MBs Read/sec	The cumulative number of host MBs read per second by the group.
Host MBs Written/sec	The cumulative number of host MBs written per second by the group.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
BE Read 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 calculated average read response time.
Host IO Limit % Time Exceeded	
Write RT (ms)	The calculated average write response time

**Table 128** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
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.
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.

**Table 128** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
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 sequential.
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$



**Table 128** Storage group metrics (continued)

Metric	Description
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the group.
BE MBs Transferred/sec	Calculated value: (MBs read per sec + MBs written per sec)
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	Total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes to make a request by the disk directors to the cache.
BE Disk Read 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	

**Table 128** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
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)	
Write RT Over 32	
Read RT Count 0 to 1	
Response Time (ms)	The calculated average response time.
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.

**Table 128** Storage group metrics (continued)

<b>Metric</b>	<b>Description</b>
BE RDF Copy	Internal metric.
BE RDF Copy (MB)	Internal metric.
Allocated Capacity (GB)	The total allocated storage group capacity.
Total Capacity (GB)	The total capacity of the storage group (GBs).
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.
Compressed Tracks	The number of compressed tracks for the storage group (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the storage group tracks that are compressed tracks (applies only to storage systems running Enginuity OS 5876).
Write paced delay	
Avg. Write Paced Delay (ms)	
Sync Write Reqs/sec	The number of sync write requests managed by this storage group per second.
Sync MBs Sent/sec	The number of sync MBs sent per second via this storage group.
Async Write Reqs/sec	The number of async write requests managed by this storage group per second.
Async MBs Sent/sec	The number of async MBs sent per second via this storage group.
Information Alert Count	Number of information alerts.
Warning Alert Count	Number of warning alerts.
Critical Alert Count	Number of critical alerts.
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Read Hits/sec	
RDF Rewrites/sec	
% Virtual Provisioning Space Saved	

**Table 128** Storage group metrics (continued)

Metric	Description
Compression Ratio	<p><b>Note</b></p> <p>Compression ratio metrics are supported only on systems capable of compression. Any value reported on a system that is not capable of compression should be disregarded.</p>
QOS Delayed IO Count	The amount of IO delayed that was introduced by enabling QoS services.
QOS Delayed IO Time ( $\mu$ s )	The amount of time IO was delayed that was introduced by enabling QoS services.
QOS Host Delayed IO Count	The amount of Host IO delayed that was introduced by enabling QoS services.
QOS Host Delayed IO Time ( $\mu$ s )	The amount of time IO was delayed that was introduced by enabling QoS services.
Response Time (ms)	The calculated average response time.
PowerPath Average Response Time (ms)	The average response time in milliseconds using PowerPath metrics.
PowerPath Observed Delta RT (ms)	The difference between PowerPath average response time and device/storage group response time.
PowerPath Observed Relative RT (%)	The relative response times between the PowerPath average response time and the device/storage group response time.
Tracks Deduped Not Written	Indicates number of LRU writes deduped and not producing backend IO. This metric will be used for rate calculation.
Writes Broke Dedup Relationships	Indicates number of writes forcing dedup-split ( $2 \times 128K$ IO's and allocation expense). This metric will be used for rate calculation.
Duplicate Track Count	Indicates capacity cost of disabling de-dup for a given device. This is used to indicate the number of de-dup track at the instance of the query and not used for rate calculation.

## Storage Group by Pool metrics

The following table lists all metrics that are available for storage groups (by pool).

**Table 129** 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.

**Table 129** Storage group (by pool) metrics (continued)

Metric	Description
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.
Allocated Tracks	
Compressed Tracks	The number of compressed tracks in the storage group (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the total tracks that are compressed (applies only to storage systems running Enginuity OS 5876).
Local Replication Tracks	

## Storage Group by Tier metrics

The following table lists all metrics that are available for storage groups (by tier).

**Table 130** 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 Tracks	The number of allocated tracks in the storage group (by tier).
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.)

**Table 130** Storage group (by tier) metrics (continued)

Metric	Description
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 (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the storage group tracks that are compressed tracks (applies only to storage systems running Enginuity OS 5876).
Local Replication Tracks	

## Thin Pool metrics

The following table lists all metrics that are available for thin pools.

**Table 131** 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.

**Table 131** Thin pool metrics (continued)

<b>Metric</b>	<b>Description</b>
WP Count	The number of tracks currently in write pending mode for the thin pool.
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)
% 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.
Ingress Tracks	The number of tracks entering the pool.
Egress Tracks	The number of tracks leaving the pool.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
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)	Internal metric.
BE XOR Reads	Internal metric.

**Table 131** Thin pool metrics (continued)

Metric	Description
BE XOR Reads (KB)	Internal metric.
BE Reads For Copy	Internal metric.
BE Reads For Copy (KB)	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Blocks Written For Copy (KB)	
BE Blocks Read For VLUN Migration	Internal metric.
BE Reads For VLUN Migration	Internal metric.
BE Writes For VLUN Migration	Internal metric.
BE Blocks Written For VLUN Migration	
BE Writes For Rebuild	Internal metric.
BE Prefetched Tracks/sec	
BE Writes For Rebuild (KB)	Internal metric.
BE RDF Copy	Internal metric.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE RDF Copy (MB)	Internal metric.
Total Capacity (GB)	The total capacity of the thin pool (GBs).
Device Block Size	The block size of the volume.
Total Tracks	The number of allocated tracks in the thin pool.
Compressed Tracks	The number of compressed tracks in the thin pool (applies only to storage systems running Enginuity OS 5876) .
Compression Scan Reads	
% Compressed Tracks	The percent of the total tracks that are compressed (applies only to storage systems running Enginuity OS 5876).

## Thin Volume metrics

The following table lists all metrics that are available for thin volumes.

**Table 132** 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.



**Table 132** Thin volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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	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 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 calculated average read response time.
Write RT (ms)	The calculated average write response time.
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.

**Table 132** Thin volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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 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.

**Table 132** Thin volume metrics (continued)

<b>Metric</b>	<b>Description</b>
Random IOs/sec	The number of IOs from a host not identified as part of a sequential stream.
Random Reads/sec	The number of read IO commands from a host not identified as part of a sequential stream.
Random Writes/sec	The number of write IO commands from a host not identified as part of a sequential stream.
Random Read Hits/sec	The number of random read IOs that were satisfied from the cache.
Random Write Hits/sec	The number of random write IOs that were immediately placed in cache because space was available.
Random Read Misses/sec	The number of random read IOs that were misses.
Random Write Misses/sec	The number of random write IOs that were misses.
Avg IO Size (KB)	Calculated value: (HA Kbytes transferred per sec / total IOs per sec)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq ios per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$

**Table 132** Thin volume metrics (continued)

<b>Metric</b>	<b>Description</b>
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the volume.
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.
% 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 cache to disks.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache in a second.
Number of RDF Invalid Tracks	
RDF Response Time (ms)	
BE RDF Copy	
BE RDF Copy (MB)	
Allocated Capacity (GB)	The allocated capacity of the thin volume (GB).
Total Capacity (GB)	The total capacity of the thin volume (GB).

**Table 132** Thin volume metrics (continued)

<b>Metric</b>	<b>Description</b>
Device Block Size	The block size of the volume.
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 (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	
Random Writes to WP slots per second	Number of writes to affect the back end.
Random Writes to WB Slots in KB per second.	The amount of writes to affect the back end (KB).
QOS Delayed IO Count	The amount of IO delayed that was introduced by enabling QoS services.
QOS Delayed IO Time ( $\mu$ s )	The amount of time IO was delayed that was introduced by enabling QoS services.
QOS Host Delayed IO Count	The amount of Host IO delayed that was introduced by enabling QoS services.
QOS Host Delayed IO Time ( $\mu$ s )	The amount of time IO was delayed that was introduced by enabling QoS services.
Response time (ms)	The calculated average response time.
Response Time (ms)	The calculated average response time.
PowerPath Average Response Time (ms)	The average response time in milliseconds using PowerPath metrics.
PowerPath Observed Delta RT (ms)	The difference between PowerPath average response time and device/storage group response time.
PowerPath Observed Relative RT (%)	The relative response times between the PowerPath average response time and the device/storage group response time.

**Table 132** Thin volume metrics (continued)

Metric	Description
Tracks Deduped Not Written	Indicates number of LRU writes deduped and not producing backend IO. This metric will be used for rate calculation.
Writes Broke Dedup Relationships	Indicates number of writes forcing dedup-split (2*128K IO's and allocation expense). This metric will be used for rate calculation.
Duplicate Track Count	Indicates capacity cost of disabling de-dup for a given device. This is used to indicate the number of de-dup track at the instance of the query and not used for rate calculation.

## Tier by Storage Group metrics

The following table lists all metrics that are available for tiers (by storage group).

**Table 133** 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	The number of writes per second that were destaged to disk.
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	The size (MBs) of writes per second that were destaged to disk.
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 Tracks	The number of allocated tracks in the storage group.

**Table 133** Tier (by storage group) metrics (continued)

Metric	Description
Compressed Tracks	The number of compressed tracks for the storage group (applies only to storage systems running Enginuity OS 5876).
% Compressed Tracks	The percent of the storage group tracks that are compressed tracks (applies only to storage systems running Enginuity OS 5876).
Local Replication Tracks	

## Virtual Pool Tier metrics

The following table lists all metrics that are available for virtual pool tiers.

**Table 134** Virtual pool tier metrics

Metric	Description
Allocated Capacity (GB)	Allocated capacity in GBs.
Allocated Pool Capacity (GB)	The allocated pool capacity in GBs.
Total Tracks	The total number of tracks for the virtual pool.
BE Disk Read RT (ms)	The average time it takes read requests from the disk directors to cache.
BE Read Request 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.
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)
BE MBs Read/sec	The number of MBs read each second by the disk directors.
BE MBs Transferred/sec	The MBs transferred per second to the disk directors.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE % Reads	The percent of the back-end IO that were read requests.
BE % Writes	The percent of the back-end IO that were write requests.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.

**Table 134** Virtual pool tier metrics (continued)

<b>Metric</b>	<b>Description</b>
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Req/sec	The number of read requests each second performed by the disk directors to cache.
BE RT (ms)	
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 Write Req/sec	The number of write requests each second performed by the disk directors to cache.
Compressed Tracks	The number of compressed tracks in the thin pool.(applies only to storage systems running Enginuity OS 5876).
BE RDF Copy (MB)	Internal metric.
BE Reads For Copy (KB)	Internal metric.
BE Blocks Read For VLUN Migration	
BE Blocks Written For Copy (KB)	
BE Writes For Rebuild (KB)	Internal metric.
BE Blocks Written For VLUN Migration	
BE Optimize Writes (KB)	Internal metric.
BE Partial Sector Writes (KB)	The number of partial sector writes by the disk directors.
BE RDF Copy	Internal metric.
BE Reads For VLUN Migration	Internal metric.
BE Reads For Copy	Internal metric.
BE Writes For Copy (KB)	Internal metric.
BE Writes For Rebuild	Internal metric.
BE Writes For VLUN Migration	Internal metric.
BE XOR Reads (KB)	Internal metric.
BE XOR Reads	Internal metric.
Device Block Size	The block size of the volume.
Total Capacity (GB)	The total capacity of the virtual pool tier.
Egress Tracks	The number of tracks leaving the pool.
Enabled Pool Capacity (GB)	The enabled pool capacity in GBs.
Ingress Tracks	The number of tracks entering the pool.



**Table 134** Virtual pool tier metrics (continued)

<b>Metric</b>	<b>Description</b>
IO Density	The number of BE requests per GB of disk.
% Compressed Tracks	The percent of the total tracks that are compressed.
% Reads	The percentage of IO operations that were reads.
% Used Capacity	The percent of the virtual pool tier that is used.
% Write	The percentage of IO operations that were writes.
BE Prefetched MBs/sec	The number of MBs prefetched from disk to cache per second.
Response Time (ms)	The average time it takes to satisfy IO requests.
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.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to cache.
Total Pool Capacity (GB)	The total capacity of the virtual pool.
Used Pool Capacity (GB)	The used pool capacity in GBs.
WP Count	The number of tracks currently in write pending mode.

## Volume metrics

The following table lists all metrics that are available for volumes.

**Table 135** Volume metrics

<b>Metric</b>	<b>Description</b>
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.

**Table 135** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
Host Hits/sec	The total number of host read IO and write IO operations performed each second by the Symmetrix volume that were immediately satisfied by cache.
Host Read Hits/sec	The total number of host read IO operations performed each second by the Symmetrix volume that were immediately satisfied by cache.
Host Write Hits/sec	The total number of host write IO operations performed each second by the Symmetrix volume that were immediately satisfied by cache.
Host Misses/sec	The total number of host read IO and write IO operations missed each second by the Symmetrix volume.
Host Read Misses/sec	The total number of host read IO operations missed each second by the Symmetrix volume.
Host Write Misses/sec	The total number of host write IO operations missed each second by the Symmetrix volume.
Host MBs/sec	A cumulative number of host MBs read/writes per second.
Host MBs Read/sec	A cumulative number of host MBs read per second.
Host MBs Written/sec	A cumulative number of host MBs written per second.
BE Reqs/sec	The number of read/write requests each second performed by the disk directors to the cache.
BE Read Reqs/sec	The number of read requests each second performed by the disk directors to the cache.
Parity Generation Kbytes Read	
Parity Generation Reads	
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.

**Table 135** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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 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.

**Table 135** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
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)
Avg Read Size (KB)	Calculated value: (Kbytes read per sec / total reads per sec)
Avg Write Size (KB)	Calculated value: (Kbytes written per sec / total writes per sec)
% Sequential IO	Calculated value: $100 * (\text{total seq IOs per sec} / \text{total IOs per sec})$
% Seq Read	Calculated value: $100 * (\text{seq reads per sec} / \text{total IOs per sec})$
% Seq Read Hit	The percent of the sequential read operations that were immediately satisfied from cache.
% Seq Read Miss	The percent of the sequential read operations that were misses.
% Seq Writes	Calculated value: $100 * (\text{seq writes per sec} / \text{total IOs per sec})$

**Table 135** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
% Seq Write Hit	The percent of the sequential write operations that were immediately satisfied from cache.
% Seq Write Miss	The percent of the sequential write operations that were misses.
% Random IO	The percent of IO operations that were random.
% Random Read Hit	Calculated value: $100 * (\text{random read hits per sec} / \text{total IOs per sec})$
% Random Read Miss	Calculated value: $100 * (\text{random read misses per sec} / \text{total IOs per sec})$
% Random Write Hit	Calculated value: $100 * (\text{random write hits per sec} / \text{total IOs per sec})$
% Random Write Miss	Calculated value: $100 * (\text{random write misses per sec} / \text{total IOs per sec})$
Max WP Threshold	The maximum number of write-pending slots available for the Symmetrix volume.
BE MBs Transferred/sec	The number of MBs read per sec + MBs written per sec.
BE MBs Read/sec	The number of MBs read by the disk directors from the disk each second.
BE MBs Written/sec	The number of MBs written to the disk from the disk director each second.
BE Prefetched Tracks/sec	The total prefetched tracks each second from the disk directors to the cache.
BE Prefetched Tracks Used/sec	The number of prefetched tracks used each second from the disk directors to the cache.
BE Read Request Time (ms)	The average time it takes 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 - 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.

**Table 135** Volume metrics (continued)

<b>Metric</b>	<b>Description</b>
% 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	The number of BE requests per GB of disk.
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)	Internal metric.
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.

**Table 135** Volume metrics (continued)

Metric	Description
Allocated Capacity (GB)	The allocated volume capacity.
Total Capacity (GB)	The total capacity of the volume.
Device Block Size	The block size of the volume.
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 (ms)	
Avg. Write Paced Delay (ms)	
RDF Reads/sec	
RDF Writes/sec	
RDF MBs Read/sec	
RDF MBs Written/sec	
RDF Rewrites/sec	

## Setting the time range for viewing data

In both the Monitor view and the Analyze view, you can configure a custom time range for which data is displayed.

To set the time range for viewing data:


### Procedure

1. In either the Monitor or Analyze view, click the custom time range button.
2. Configure a start date and a start time for the time range.
3. Configure an end date and an end time for the time range.
4. Click **OK**.

## Creating a template dashboard from an Analyze view

To create a dashboard template from an Analyze view:

### Procedure

1. Select component(s) from the table in an Analyze view.
2. Click  > **Navigate to Charts**  
The selected components are displayed in a **New Chart** view.
3. Select the metrics to be displayed.

4. Click **Create**.

5.

Click  > **Save as Template**.

The **Save as Template** dialog box displays.

6. Specify a template name.

7. Click **OK** to save the template.



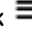
## Filtering performance data

Unisphere includes different ways to filter performance data depending on the situation:

- Filter heatmaps by context. This allows you to filter according to an instance of a particular object. For example, a particular storage group or database.
- Filter arrays and metrics by name when creating charts.
- Filter analyze views, reports, and traces by any of the headings in the object list.


## Filtering object lists

### Procedure

1. Click the filter icon .
2. Do one of the following:
  - In any of the fields displayed, type all or part of the value for which you want to filter.
  - To change the fields displayed, click  and select the **Columns** to display for filtering purposes.
  - To clear the filtering criteria, click  and select one of the following options:
    - **Reset to Defaults**
    - **Clear all filters**

## Filtering heatmaps

### Procedure

1. Click the filter icon  to open the filter dialog box.
2. From the **Context** menu, select an object type. The instances list updates to show the objects of that type.
3. Select an object from the instances list. You can filter the instances list by typing all or part of the object's name in the search field.
4. Click **OK**.



# CHAPTER 8

## Database Storage Analyzer

Topics include:

• <a href="#">Introduction</a> .....	834
• <a href="#">Database collection and retention policy</a> .....	834
• <a href="#">Mapping files</a> .....	835
• <a href="#">Viewing Databases page</a> .....	835
• <a href="#">Viewing Database Administration page</a> .....	837
• <a href="#">Registering a monitored environment</a> .....	839
• <a href="#">Adding an Oracle database</a> .....	839
• <a href="#">Registering Monitored Environment - Advanced Options</a> .....	840
• <a href="#">Adding monitored MS SQL server instances</a> .....	841
• <a href="#">Editing monitored Oracle databases</a> .....	843
• <a href="#">Editing monitored MS SQL server instances</a> .....	844
• <a href="#">Starting statistics collection</a> .....	845
• <a href="#">Stopping statistics collection</a> .....	845
• <a href="#">Running device mapping</a> .....	845
• <a href="#">Schedule device mapping</a> .....	845
• <a href="#">Removing monitored environment instance</a> .....	846
• <a href="#">Viewing the Performance Page</a> .....	846
• <a href="#">Viewing the Analytics Page</a> .....	849
• <a href="#">Viewing analytics details</a> .....	853
• <a href="#">Viewing database storage details</a> .....	853
• <a href="#">Viewing database details</a> .....	854
• <a href="#">Adding hints</a> .....	854
• <a href="#">Viewing hints</a> .....	855
• <a href="#">Editing hints</a> .....	856
• <a href="#">Enabling hints</a> .....	857
• <a href="#">Disabling hints</a> .....	857
• <a href="#">Removing hints</a> .....	858
• <a href="#">Viewing hint logs</a> .....	858
• <a href="#">Hint operations</a> .....	859

## Introduction

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

DSA is accessible by any Unisphere user. A DSA only user can also be created that has read only Unisphere access and can only view Databases section and Performance. Details on how to create users can be found at [Creating local users](#) on page 76.

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

DSA offers the following benefits:

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

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

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

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

---

### Note

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

---

## Database collection and retention policy

### For Oracle

DSA collects information by connecting directly to the monitored database through a database user. This read-only user only has select permissions on a fixed list of Oracle dictionary tables. DSA fetches data every five minutes and sends it back to the

Unisphere repository database where it aggregates the data into hourly and daily aggregations. By default, DSA saves the fetched data for fifteen days; however, this can be extended to thirty days. DSA saves the hourly aggregations for fifteen months and the daily aggregations for two years. To connect to the monitored database, verify that the database TNS port (usually 1521 or 1525) between the Unisphere repository server and the monitored database server is open.

#### For MS SQL Server

DSA collects information by connecting directly to the monitored database through a database user. This read-only user only has select permissions on a fixed list of MS SQL Server dictionary tables. DSA fetches the data directly from the MS SQL Server tables every second. The read only user can be created during the installation process or DBAs can create it manually prior to running the “Add database” option. The DSA user needs to have SYSADMIN privileges.

Note: a SYSADMIN role is required in order to collect object data. However, DSA login can be created with read only privileges, yet any functionality associated with object collection is disabled. DSA works with either MS SQL Server local user or as a Windows authenticated user.

## Mapping files

The mapping process is responsible for mapping the Oracle and MS SQL Server files to the storage system devices. By default, the process runs once a week, however, it can be configured to run at different times.

An executable called Mapper is copied to the remote database server with a list of Oracle or MS SQL Server data files to map. This process identifies the host physical devices associated with the Oracle or MS SQL Server files, and then sends the list of storage devices back to be loaded into the Unisphere DSA repository.

For Oracle databases running on Linux or Unix OS, DSA requires a read only user on the database host that has sudo privileges to run the mapper script. A user is root or full root privileges can also be used. For MS SQL Server databases, the user needs to install a DSA Listener on the monitored environment. The listener should be manually installed before adding the new monitored database. In order to install the DSA listener on the monitored environment, go to the administration tab in DSA, click **Register** and then select the MS SQL Server option. Click DSA Listener to download it, copy it to the monitored environment and then follow the instructions in the readme.txt file. Note that the default port for the DSA listener is 22200. This is configurable from the DSA client. DSA mapper should not impact database activity on the monitored server.

In order to run the mapping on a database that is running on VMware virtual disks, make sure to add the virtual server to Unisphere. On the Unisphere client, go to **HOME** and then select **VMWARE**. Click **Register vCenter/ESXi** to register an ESXi server. In addition, set the parameter `disk.EnableUUID = "TRUE"` in the ESXi server.

Virtual environments other than VMware are not supported, for example Solaris containers, AIX VIO and VPLEX.

## Viewing Databases page

This view lists the monitored databases that have been registered. Both MS SQL Server and Oracle are displayed on the same list. This view allows you to quickly

identity a potential problem with a database by comparing the storage and database KPIs.




For an introduction to Database Storage Analyzer, see [Introduction](#) on page 834

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Databases** to view the full list of monitored environments and their associated properties.


The following properties display for each monitored environment:

- **Environment Name**—Name of the monitored environment.
- **DB Type**— Type of database or instance.
- **DB RT Status**— Percentage of occurrences where the response time was as expected, over, poor or missing samples. Hovering over the DB RT Status gives a breakdown of the percentage of occurrences where the database read response time was as expected, over, poor or missing samples.
- **Status**— Status of database that takes into account collections, mapper, synchronizer and statistics status.
- **DB Read RT (ms)**—Average database read response time in milliseconds.
- **Storage Read RT (ms)**—Average storage read response time in milliseconds.
- **DB IOPs**—Average Database Input/Output (I/O) per second (IOPS) is the number of reads from data file and temporary files and the number of writes to data files, temporary files, redo logs and archive files per second.
- **Storage IOPs**— Average Storage IOPS is the number of reads and writes per second to the storage volumes that are associated with the database files.

3. Optional: Click and drag the properties columns to rearrange the column order.
4. Optional: Click the drop down menu to select the time range for the monitoring activity.
5. Optional: Create a filter for the list by clicking  and typing any combination of properties.
6. Optional: Click  to reset the filters to defaults setting, clear all filters or change the displayed columns.
7. Optional: Click  to open the details panel. Click an environment name instance to view summary Database, Storage and Statistics information about the environment.

The following properties display for each monitored environment:

- **Database**—the following properties display:
  - **Environment Name**— Name of the monitored environment.
  - **DB Host**— Hostname or IP address of the server the database is running on.
  - **DB Type**—Type of database or instance.
  - **DB Version**— Version of database software the database is running on the monitored environment.

- **Database Size (GB)**— Size of database.
  - **DB Collection status**—Status of database collection. Clicking on the status icon brings you to the Administration tab where information about the status is displayed.
  - **DB Mapping Status**—Status of database mapping. Clicking on the status icon brings you to the Administration tab where information about the status is displayed.
  - **Storage Collection Status**—Status of storage collection. Clicking on the status icon brings you to the Administration tab where information about the status is displayed.
  - **Storage**—the following properties display:
    - **Array**— Identity of the storage system that the database or instance is running on.
    - **Storage Groups**— Storage Groups associated with the database. Click the number next to **Storage Groups** to go to the Storage Section.
    - **SG Compliance**— Status indicating if the associated Storage Groups are in compliance with its service level.
    - **Front End**— Displays the worst performing FE port or director port busy threshold status.
    - **Capacity (GB)**— Database capacity in Gigabytes.
    - **Number of Volumes**— Number of volumes in storage groups.
    - **Masking Views**— Number of masking views.
  - **Statistics**—the following properties display:
    - **DB Redo RT (ms)**—Average database Redo Write response time for the selected period of time based on Oracle statistics.
    - **Storage Redo Write RT (ms)**—Average storage write response time for storage volumes where the database redo files are located.
    - **DB IOPs Reads**— Average instance Input/Output (I/O) per second (IOPS) reads.
    - **DB IOPs Writes**— Average instance Input/Output (I/O) per second (IOPS) writes.
    - **DB IOPs Log Writes**—Average instance Input/Output (I/O) per second (IOPS) log writes.
    - **DB Reads (MB/s)**— Average DB reads in milliseconds.
    - **DB Writes (MB/s)**— Average DB writes from in milliseconds.
    - **DB Log Writes (MB/s)**— Average DB log writes from in milliseconds.
    - **DB Instance R/W Ratio**— Read to Write ratio for the database.
8. Optional: Click  to open the details panel. Click **VIEW ALL DETAILS** to go to the **Performance** section for the monitored environment.
9. Optional: Double click an environment name instance to go the **Performance** section for the monitored environment.

## Viewing Database Administration page

This task allows you to view the environment instances being monitored by DSA.

---

**Note**

To ensure that storage performance statistics are collected you need to enable Performance collection in Unisphere (see [Registering storage systems](#) on page 592) .

---


For an introduction to Database Storage Analyzer, see [Introduction](#) on page 834


**Procedure**

1. Click **HOME**.
2. Select **DATABASES > Administration** to view the full list of monitored environments and their associated properties.

The following properties display for each monitored environment:


- **DB Name**—Name or IP address of the monitored environment or instance host.
- **DB Type**—Type of database.
- **DB Version**—monitored environment or instance version.
- **Array**—Identity of the storage system that the database or instance is running on.
- **Status**—Icon indicating the status of the system.

3. Optional: Create a filter for the list by clicking  and defining any combination of Environment Name, Database Host, Database Type, Database Version, Array, System Status and Mapper Status.

4. Optional: Click  to reset the filters to defaults setting, clear all filters or change the displayed columns.

The following controls are available:

- **Register**—[Registering a monitored environment](#) on page 839
- **Edit**—[Editing monitored Oracle databases](#) on page 843 or [Editing monitored MS SQL server instances](#) on page 844
- **Run Mapping**—[Running device mapping](#) on page 845
- **Start Collection**—[Starting statistics collection](#) on page 845
- **Stop Collection**—[Stopping statistics collection](#) on page 845
- **Unregister**—[Removing monitored environment instance](#) on page 846
- **Schedule Mapping Configuration**—[Schedule device mapping](#) on page 845

5. Optional: Select an environment instance and click  to view details about the environment.

The following properties display:

- **Environment Name**—Name of the monitored environment.
- **Last Mapping Date**—Date of last mapping of environment.
- **Installation Date**—Date the environment was installed.
- **Oldest Date**—Oldest date.

- **Latest Available Data**—Date of latest available data.
- **Mapper Host**—Host where the mapping process is executed.
- **Mapper State**—Pass or Fail status of mapping process.
- **DB Collection state**—Database collection state.
- **DB Collection status**—Status of database collection.
- **DB Mapping Status**—Status of database mapping.
- **Storage Collection Status**—Status of storage collection.

## Registering a monitored environment

- Registering an Oracle database (see [Adding an Oracle database](#) on page 839).
- Registering a MS SQL server (see [Adding monitored MS SQL server instances](#) on page 841).

## Adding an Oracle database

This task allows you to add an Oracle database to be monitored by DSA. For discovery of a new database a user on the database will be required to be given read only GRANTS to certain v\$ tables. Details of the GRANTS script can be found at (<Unisphere server IP address>/univmax/database/Oracle/create\_emc\_unisphere\_guest\_user.sql ). A user on the database host needs to be supplied with sudo access to run the mapper script. Details of what to update in the sudoers files are supplied on the mapping step of the wizard.

### Procedure

1. Click **HOME**, then click **DATABASES** and then click **Administration** to view the full list of monitored environments and their associated attributes.
2. Click **Register**.
3. Select **Oracle** and click **NEXT**.
4. Enter values for the following parameters:
  - **DB Host Name or IP Address**—Name or IP address of the database host.
  - **Service**—Service name that may contain up to 16 alpha numeric characters.
  - **SID**—System identifier (SID) is a unique name for an Oracle database instance on a specific host.
  - **Database Port**—Database listener port number.
5. Click **NEXT**.
6. Select one of the following:
  - **Create DB User**—Creating a DSA user is used when a new DSA database user is added by the installation process. The sys user must be provided in SYS Password.
  - **Existing DB User**—Using an existing DSA user requires creating the user manually prior to the installation using the script provided.
7. Enter values for the following parameters:

- **DB Username**—Name of the database user.
- **DB Password**—User password for the database.

If you select the **Create DB User** option to create a new DSA user, the user can only connect to the database via the MS SQL server authentication mode. If you select **Existing DB Login**, the user also has the option to connect using Windows authentication mode.

Creating a DSA user is used when a new DSA login is added by the installation process. A SYSADMIN login must be provided during the next step of the installation.

Using an existing DSA login requires creating the user manually prior to the installation by running Create User Script (click **Create User Script** to access the script) provided using the SYSADMIN login, and then manually entering the new user during the installation. Note that SYSADMIN role is required in order to collect object data. However, DSA login can be created with read only privileges, yet any functionality associated with object collection is disabled. In order to map database files to Symmetrix devices, DSA needs to install a DSA Listener on the monitored environment. The listener can be installed by the installation process or manually. In order to install the DSA listener manually on the monitored environment, click **DSA Listener** to download it, copy it to the monitored environment host, extract to `C:\EmcDBAgent` and follow the instructions in the readme.txt file.

8. Click **NEXT**.

9. For non-root users, enter values for the following parameters:

- **Mapper Host Name or IP Address**—Host name or IP address where the mapping process is executed.
- **SSH Port**—Port where the mapping process is executed.
- **SSH Username**—SSH username.

---

#### Note

SSH needs to be configured on your system and you need to confirm that the appropriate permissions are set for the sudo login user in the sudoers file.

---

- **Login User**—Sudo user name (mandatory field).
- **Login Password**—Sudo user password (mandatory field).
- **SSH Password**—SSH password.

10. Click **NEXT**.

11. Click **Advanced Options** (see [Registering Monitored Environment - Advanced Options](#) on page 840)

12. Click **NEXT**.

13. Review details and click **FINISH**.

## Registering Monitored Environment - Advanced Options

### Procedure

1. Oracle only: Enter or select values for the following parameters :



- **Use <login user id> home directory to store mapping files**— Select this option to use the login user's home directory to store the mapping files. The temporary mapping files are stored by default in temp directories. Copy the mapper command line information by clicking on the "i" button and edit the user entry in the sudoers file. for example: `test ALL=NOPASSWD:/home/test/UnisphereMapper_hostname_SID_on_*.*/mapper_agent.sh map -file dbfiles.dat -log_file dbc_mapper.log`. Ensure that the user has permissions for the /home/user directory.
  - **Alternative path to sudo**—Use an alternative path to sudo if sudo is not available by default or you would like to specify a different sudo path.
  - **Run sudo interactively**— Select this option to run substitute user do (sudo) interactively. This allows users to run **Use login user home directory to store mapping files** programs with the security privileges of another user (normally the superuser, or root).
  - **Disable mapping extents**— Select this option for extremely large databases, when the you wish to monitor a large number of databases, or if there are issues mapping a database. With mapping extents disabled, database objects with not be mapped to storage devices. As a result, the Database Objects list view on the Analytics Tab will not contain database objects. Hinting will be disabled but all other functionality will work as normal.
2. MS SQL server only: Select a value for the following parameter:
- Disable mapping extents**— Select this option for extremely large databases, when the you wish to monitor a large number of databases, or if there are issues mapping a database. With mapping extents disabled, database objects with not be mapped to storage devices. As a result, the Database Objects list view on the Analytics Tab will not contain database objects. Hinting will be disabled but all other functionality will work as normal.
- It is recommended to disable mapping extents if you have a MS SQL Server instance with a large number of databases.
3. Click OK.

## Adding monitored MS SQL server instances

This task allows you to add MS SQL server instances to be monitored by DSA. For discovery of a new database a user on the database will be required to be given read only GRANTS to certain tables. Details of the GRANTS script can be found at (<Unisphere server IP address>/univmax/database/ MSSQLServer/create\_emc\_unisphere\_guest\_user.sql ).

### Procedure

1. Click **HOME**, then **DATABASES** and then **Administration** to view the full list of monitored environments and their associated attributes.
2. Click **Register**.
3. Select **MS SQL Server** and click **NEXT**.
4. Enter values for the following parameters:
  - **DB Host Name or IP Address**—Name or IP address of the database host.

- **Database Name**—Name of database.
  - **Instance**—Instance name or ID.
  - **Database Port**—Database listener port number.
5. Click **NEXT**.
  6. Enter values for the following parameters:
    - **Create DB User**—Create a new database user.
    - **Existing DB User**—Login as an existing database user.
    - Click **Learn More about creating a DB User** to learn more about creating a DB user or using an existing DB user.
    - **Windows Authentication Mode**—Windows Authentication Mode or SQL Server Authentication Mode.
    - **SQL Server Authentication Mode**—SQL Server Authentication Mode.
    - **DB Username**—Database username. This is the user running the mapping executable. This user needs to be part of the Administrator group on the database server.
    - **DB Password**—Database Password. This is the windows password of the user running the mapping executable.
    - **Sys Admin Login**—SYSADMIN login name. A SYSADMIN login is temporarily used to set up the DSA login.
    - **SYS Password**—SYSADMIN password.

If you select the **Create DB User** option to create a new DSA user, the user can only connect to the database via the SQL server authentication mode. If you select **Existing DB Login**, the user also has the option to connect using Windows authentication mode.

Creating a DSA user is used when a new DSA login is added by the installation process. A SYSADMIN login must be provided during the next step of the installation.

Using an existing DSA login requires creating the user manually prior to the installation by running Create User Script (click **Create User Script** to access the script) provided using the SYSADMIN login, and then manually entering the new user during the installation. Note that SYSADMIN role is required in order to collect object data. However, DSA login can be created with read only privileges, yet any functionality associated with object collection is disabled. In order to map database files to Symmetrix devices, DSA needs to install a DSA Listener on the monitored environment. The listener can be installed by the installation process or manually. In order to install the DSA listener manually on the monitored environment, click **DSA Listener** to download it, copy it to the monitored environment host, extract to `C:\EmcDBAgent` and follow the instructions in the readme.txt file.

7. Click **NEXT**.
8. Modify any number of the following values:
  - **Mapper Host Name or IP Address**—Host name or IP address where the mapping process is executed.
  - **SSH Port**—Port where the mapping process is executed.
  - **SSH Username**—SSH username.

- **SSH Password**—Secure SSH password.
9. Click **NEXT**.
  10. Review details and click **FINISH**.

## Editing monitored Oracle databases

This task allows you to edit Oracle database records already being monitored by the Database Storage Analyzer (DSA) application.

### Procedure

1. Click **HOME**, then **DATABASES** and then **Administration** tab to view the full list of monitored environments and their associated attributes.
2. Select an Oracle database and click **Edit**.
3. Modify any number of the following values:
  - **DB Host Name or IP Address**—Name or IP address of the database host.
  - **DB Username**—Database username. This is the user running the mapping executable. This user needs to be part of the Administrator group on the database server.
  - **Service**—Service name that may contain up to 16 alpha numeric characters.
  - **SID**—System identifier (SID) is a unique name for an Oracle database instance on a specific host.
  - **DB Password**—Database Password. This is the windows password of the user running the mapping executable.
  - **Database Port**—Database listener port number.
4. Click **NEXT**.
5. Modify any number of the following values:
  - **Mapper Host Name or IP Address**—Host name or IP address where the mapping process is executed.
  - **SSH Port**—Port where the mapping process is executed.
  - **SSH Username**—SSH username.
  - **SSH Password**—Secure SSH password.

---

### Note

SSH needs to be configured on your system and you need to confirm that the appropriate permissions are set for the sudo login user in the sudoers file.

---

6. Click **NEXT**.
7. Modify any number of the following values:
  - **Use home directory to store mapping files**—Select this option to use the login user's home directory to store the mapping files. The temporary mapping files are stored by default in temp directories.
  - **Sudo access for account**—Select this option if sudo access needed for the account.
  - **Alternative path to sudo**—Use an alternative path to sudo if sudo is not available by default or you would like to specify a different sudo path.

- **Run sudo interactively**—Select this option to run substitute user do (sudo) interactively. This allows users to run programs with the security privileges of another user (normally the superuser, or root).
8. Click **NEXT**.
  9. Review details and click **FINISH**.

## Editing monitored MS SQL server instances

This task allows you to edit MS SQL server instances already being monitored by the Database Storage Analyzer (DSA) application.

### Procedure

1. Click **HOME**, then **DATABASES** and then **Administration** to view the full list of monitored environments and their associated attributes.
2. Select an instance of MS SQL Server and click **Edit** to open a dialog.
3. Modify any number of the following values:
  - **DB Host Name or IP Address**—Name or IP address of the database host.
  - **Database Name**—Name of database.
  - **Instance**—Instance or ID.
  - **Database Port**—Database listener port number.
4. Click **NEXT**.
5. Modify any number of the following values:
  - **Windows Authentication Mode**—Windows Authentication Mode or SQL Server Authentication Mode.
  - **SQL Server Authentication Mode**—SQL Server Authentication Mode.
  - **DB Username**—Database username. This is the user running the mapping executable. This user needs to be part of the Administrator group on the database server.
  - **DB Password**—Database Password. This is the windows password of the user running the mapping executable.
6. Click **NEXT**.
7. Modify any number of the following values:
  - **Mapper Host Name or IP Address**—Host name or IP address where the mapping process is executed.
  - **SSH Port**—Port where the mapping process is executed.
  - **SSH Username**—SSH username.
  - **SSH Password**—Secure SSH password.
8. Click **NEXT**.
9. Review details and click **FINISH**.

## Starting statistics collection

This task allows you to start statistics collection for the monitored environment instance.

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Administration** to view the full list of monitored environments and their associated properties.
3. Select an environment instance and click **Start Collection** to start statistics collection for the database or instance.
4. Click **OK** to confirm.

## Stopping statistics collection

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Administration** to view the list of monitored environments and their associated attributes.
3. Select an environment instance and click **Stop Collection** to stop statistics collection for the database or instance.
4. Click **OK** to confirm.

## Running device mapping

This task allows you to manually run device mapping for a monitored database or instance. During device mapping, the list of database or instance files is copied using ssh to the monitored database or instance host. A process executing on the monitored database or instance host identifies the list of host physical devices associated with those files. The list is sent back and loaded into the DSA repository.

### Procedure



1. Click **HOME**.
2. Select **DATABASES > Administration** to view the full list of monitored environments and their associated properties.
3. Select an environment and click **Run Mapping** to run device mapping for the database record.
4. Click **OK** to confirm.

## Schedule device mapping

This task allows you to schedule device mapping for a monitored database or instance. During device mapping, the list of database or instance files is copied using ssh to the monitored database or instance host. A process executing on the monitored database or instance host identifies the list of host physical devices associated with those files. The list is sent back and loaded into the DSA repository.

**Procedure**

1. Click **HOME**.
2. Select **DATABASES > Administration** to view the full list of monitored environments and their associated properties.
3.
 




 Select an environment, click  and then click **Schedule Mapping Configuration**.
4. Select day(s) of the week and the time.
5. Click **OK**.

## Removing monitored environment instance

This task allows you to remove environment instances so that they are no longer monitored by DSA.

**Procedure**


1. Click **HOME**.
2. Select **DATABASES > Administration** to view the full list of monitored environments and their associated properties.
3.
 


 Select an environment instance and click  and then click **Unregister** to remove the selected database or instance.
4. Click **OK** to confirm.

## Viewing the Performance Page

The Performance page provides database and storage information in a single view allowing Database Administrators and Storage Administrators to troubleshoot performance issues in a more timely manner.

**Procedure**

1. Click **HOME**.
2. Select **DATABASES > Databases** to view the full list of monitored environments and their associated properties.
3. Do one of the following:
  - Double click on a database name instance to navigate to the view all details section. The performance tab is displayed.
  - Click on a database name instance and click  to open the details panel. Click **VIEW ALL DETAILS** to go to the **Performance** section for the monitored environment.
4. Optional: Click the **Analytics** tab to navigate to the **Analytics** area where you can analyze the I/O characteristics of specific business processes running in the database.
5. Optional: Click the **Storage** tab to navigate to the **Storage** area where you can view database related information.

6. Optional: Click the **Hinting** tab to navigate to the **Hinting** area where you can view hint related information.
7. Optional: For MS SQL Server databases only. Click the **Database** tab to navigate to the **Databases** area where you can view the list of monitored environments.
8. Optional: On the **Performance** tab, click the **Dashboard** button to navigate to the **PERFORMANCE** area where you can view the **Performance** details for each database.
9. Optional: On the **Performance** tab, click the **Heatmap** button to view a Heatmap related to the storage associated with the selected monitored environment.
10. Click the drop-down menu to select the time range for the monitoring activity.
11. For Oracle environments, view the performance chart for **I/O Wait minutes**. The x-axis represents time and the y-axis represents the following:
  - **I/O Wait**—Time spent waiting for I/O operations to complete.
  - **Non-I/O Wait**—Time spent waiting for non-I/O operations to complete.
  - **Redo**—Time spent waiting for redo operations to complete.
  - **Background**—Time spent by background process waiting for a database resource.
12. For pre-Oracle 12 environments, view the performance chart for **Average Active Session Wait**. The x-axis represents time and the y-axis represents the accumulative time the sessions in the database are waiting for each event/class. The Oracle wait classes are listed below:
  - **Administrative**—Waits resulting from DBA commands (for example, an index rebuild).
  - **Application**—Waits resulting from user application code (for example, lock waits caused by row level locking or explicit lock commands).
  - **Cluster**—Waits related to Real Application Cluster resources (for example, global cache resources such as 'gc cr block busy').
  - **Commit**—This wait class only comprises one wait event - wait for redo log write confirmation after a commit (that is, 'log file sync').
  - **Concurrency**—Waits for internal database resources (for example, latches).
  - **Configuration**—Waits caused by inadequate configuration of database or instance resources (for example, undersized log file sizes, shared pool size).
  - **Network**—Waits related to network messaging (for example, 'SQL\*Net more data to dblink').
  - **ORA CPU**—Sessions that are consuming CPU.
  - **ORA Wait CPU**—Waits for CPU – Sessions waiting in the CPU queue.
  - **OS CPU**—Displays the non-database process waiting/consuming CPU on the host.
  - **Other**—Waits which should not typically occur on a system (for example, 'wait for EMON to spawn').
  - **Scheduler**—Resource Manager related waits (for example, 'resmgr: become active').

- **System I/O**—Waits for background process I/O (for example, DBWR wait for 'db file parallel write').
- **User I/O**—Waits for user I/O (for example 'db file sequential read').

13. For MS SQL server instances, view the performance chart for **I/O Wait vs. Non-I/O Wait**.

14. For MS SQL server instances, view the performance chart for **Wait Classes over Time**.

There are filters for the response time, IOPs, throughput and I/O size kb charts referred to in the next steps. You can turn on or off database reads, writes and redo writes (Oracle), Log Writes (MS SQL Server) and storage reads, writes and redo writes (Oracle) Log Writes (MS SQL Server) for all four charts at the same time. For MS SQL Server environments, a filter icon is visible that allows the four charts listed above to be filtered to a database rather than at the instance level.

15. View the performance chart for **Response Time**.

The x-axis represents time and the y-axis represents the response time for Reads, Writes and Redo Writes (Oracle only) or Log Writes (SQL server only) for both Database and Storage. All of the Storage statistics represents the storage devices that are associated with the database rather than the entire array. Clicking DB Read Response Time Thresholds enables the display of threshold lines on the chart. The thresholds are based on the database thresholds as defined in the system. If you hover over the line, you will see the threshold information. The first threshold for database response time is 20ms. The second threshold for database response time is 30ms.

The DB Read Response Time Thresholds can be turned on or off for the response time chart. (Thresholds are configured per environment - see [Managing Database Storage Analyzer \(DSA\) environment preferences](#) on page 85) .

There are filters ( EFD, 15k, 10k, 7.2k and DX) that can turned on or off on IOPS and throughput MB/s charts.

The charts can be filtered to show reads or writes using the drop down.

16. View the performance chart for **IOPS**.

The x-axis represents time and the y-axis represents IOPS for Reads, Writes and Redo Writes for both Database and Storage. Note that for Oracle12 PDB, the redo statistics can be seen at the CDB level only.

A IOPS Density check box will display IOPS density in the IOPS chart

17. View the performance chart for **Throughput**. The x-axis represents time and the y-axis represents Throughput for Reads, Writes and Redo Writes for both Database and Storage.

18. View the performance chart for **I/O Size kb**. The x-axis represents time and the y-axis represents I/O size for reads, writes and redo writes for both database and storage. Information about the average IO size for reads, writes and redo writes (for both database and storage).

19. View the performance chart for **IOPS** for the Storage Back-end Activity and Tier Capacity. The x-axis represents time and the y-axis represents IOPS for Reads or Writes (displayed with or without Density) for EFD, FC or SATA drives.




20. View the performance chart for **Throughput** for the Storage Back-end Activity and Tier Capacity. The x-axis represents time and the y-axis represents Throughput for Reads or Writes (displayed with or without Density) for EFD, FC or SATA drives.
21. View the performance chart for **Tier Capacity** for the Storage Back-end Activity and Tier Capacity. The x-axis represents time and the y-axis represents Tier Capacity for Reads or Writes (displayed with or without Density) for EFD, FC or SATA drives.

## Viewing the Analytics Page

The Analytics page provides a means to analyze the I/O characteristics of specific business processes running in the database. Oracle wait events represent the various event types that Oracle sessions are waiting for. This information is helpful in the analysis of database bottlenecks. Wait classes are a way of grouping events to logical groups for high level analysis. There is a filter icon on the top of the Analytics view that filters the chart and the list below it.

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Databases** to view the full list of monitored environments and their associated properties.
3. Do one of the following:
  - Double click on a database name instance to navigate to the view all details section. The performance tab is displayed.
  - Click on a database name instance and click  to open the details panel. Click **VIEW ALL DETAILS** to go to the **Performance** section for the monitored environment.
4. Click the **Analytics** tab.
5. Click the drop-down menu to select the time range for the monitoring activity.
6. For an Oracle database, click the drop down menu and select the performance chart for **IO Wait Classes**. The x-axis represents time and the y-axis represents the following wait classes:
  - **Single Block Read**—Wait class associated with Single Block read operations.
  - **Multi-Block Read**—Wait class associated with Multi-Block read operations.
  - **Direct I/O**—Wait Class that represents an I/O operation where the data is asynchronously read from the database files. It comprises of the following wait events: direct path read temp, direct path write temp, direct path write, and direct path read.
  - **Other I/O**—Wait Class that comprises of I/O events where the type of I/O (for example, random or sequential scan) is unknown.
  - **System I/O**—Wait Class associated with system I/O operations.
  - **Commit**—Wait Class associated with database commit.
7. For an Oracle database, click the drop down menu and select the performance chart for **IO Wait Events**. The x-axis represents time and the y-axis represents the following wait classes:

- **Application**—Waits resulting from user application code (for example, lock waits caused by row level locking or explicit lock commands).
  - **Commit**—This wait class only comprises one wait event - wait for redo log write confirmation after a commit (that is, 'log file sync').
  - **Concurrency**—Waits for internal database resources (for example, latches).
  - **Configuration**—Waits caused by inadequate configuration of database or instance resources (for example, undersized log file sizes, shared pool size).
  - **CPU + Wait for CPU**—Sessions that are consuming CPU or waiting in the CPU queue.
  - **System I/O**—Waits for background process I/O (for example, DBWR wait for 'db file parallel write').
  - **User I/O**—Waits for user I/O (for example 'db file sequential read').
  - **Other**—Waits which should not typically occur on a system (for example, 'wait for EMON to spawn').
  - **Direct I/O**—Wait class that represents an I/O operation where the data is asynchronously read from the database files. It comprises of the following wait events: direct path read temp, direct path write temp, direct path write, and direct path read.
  - **Single Block I/O**—Wait class associated with Single Block I/O operations.
  - **Commit**—Wait class associated with database commit.
  - **Multi-Block I/O**—Wait class associated with Multi-Block I/O operations.
  - **Other I/O**—Wait class that comprises of I/O events where the type of I/O (for example, random or sequential scan) is unknown.
8. For an Oracle database, click the drop down menu and select the performance chart for **All Wait Classes**. Wait events represent the exact wait type the Oracle process is waiting on. Every wait event belongs to a class of wait event. The x-axis represents time and the y-axis represents the accumulative time the sessions in the database are waiting for each event/class.
  9. For an Oracle database, click the drop down menu and select the performance chart for **All Wait Events**. Wait events represent the exact wait type the Oracle process is waiting on. Every wait event belongs to a class of wait event. The x-axis represents time and the y-axis represents the accumulative time the sessions in the database are waiting for each event/class.
  10. For a MS SQL server instance, select the **Wait Classes** context and view the performance chart for **Analytics Wait Classes**. The x-axis represents time and the y-axis represents the accumulative time the sessions in the database where waiting for each wait event/class. The SQL server wait classes are listed below:
    - **Buffer I/O**
    - **Buffer Latch**
    - **CPU**
    - **Latch**
    - **Locks**
    - **Logging**
    - **Memory**
    - **Network I/O**

- **Waiting for CPU**
  - **Other**
11. For a MS SQL server instance, select the **Wait Events** context , and view the performance chart for **Analytics Wait Events**. Wait events represent the exact wait type the MS SQL server process is waiting on. Every wait event belongs to a class of wait event. The x-axis represents the date and the y-axis represents the accumulative time the sessions in the database where waiting for each wait event/class.
- **Running**
  - **PAGEIOLATCH\_SH**
  - **Runnable**
  - **LCK\_M\_U**
  - **LATCH\_EX**
  - **Null event**
  - **LCK\_M\_IX**
  - **PREEMPTIVE\_OS\_DECRYPTMESSAGE**
  - **LCK\_M\_X**
  - **OTHERS**
  - **PREEMPTIVE\_OS\_ENCRYPTMESSAGE**
12. View the properties of the Oracle **Database Objects**.
- **Owner**—Owner of object, for example, SYSTEM.
  - **Object Name**—In the case of a partition object, the partition name is displayed.
  - **Sub-Object Name**—In the case where the object is partitioned, the sub-object name contains the partition name.
  - **Object Type**—Type of object, for example, table or index.
  - **Object Size**—Size of the object in MB.
  - **IO Wait Time**—I/O Wait time in seconds.
  - **IO Type**—I/O Type is one of the following:
    - Single Block Read
    - Multi Block Read
    - Direct I/O
    - System I/O
    - Commit I/O
    - Other I/O
    - Row I/O
  - **Hints**—Database Hints.
13. View the properties of the MS SQL server instance objects.
- **Database name**—Name of the database.
  - **Schema name**—Name of the schema.
  - **Object Name**—In the case of a partition object, the partition name is displayed.

- **Partition**—Partition name.
- **Type**—Type of object, for example, table or index.
- **Object Size**—Size of the object in MB.
- **File Group**—Associated file group.
- **IO Wait Time**—I/O Wait time in seconds.

14. View the properties of the Oracle **Array Devices**.

- **Host Device**—Host devices that are associated with the monitored environment.
- **Array Device**— storage system devices associated with the monitored environment.
- **IO Wait Time**—I/O Wait time in seconds.
- **IO Type**—I/O Type is one of the following:
  - Single Block Read
  - Multi Block Read
  - Direct I/O
  - System I/O
  - Commit I/O
  - Other I/O
  - Row I/O

You can also view Oracle database session contextual information. The following is the list of session context types:

- Program Names
- Login Users
- Machines
- Modules
- Services
- Nodes

You can also view MS SQL server database session contextual information. The following is the list of session context types:


- Hosts
- Program Names
- NT Domains
- NT Users
- Login Names
- Original Login Names
- Login Databases

The following controls are available:

- **Details**—[Viewing analytics details](#) on page 853
- **Add to Hint**—[Adding hints](#) on page 854

## Viewing analytics details

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Databases** to view the full list of monitored environments and their associated properties.
3. Do one of the following:
  - Double click on a database name instance to navigate to the view all details section. The performance tab is displayed.
  - Click on a database name instance and click  to open the details panel. Click **VIEW ALL DETAILS** to go to the **Performance** section for the monitored environment.
4. Click the drop-down menu to select the time range for the monitoring activity.
5. Click the **Analytics** tab.

You can also view Oracle database session contextual information. The following is the list of session context types:

- Program Names
- Login Users
- Machines
- Modules
- Services
- Nodes

You can also view MS SQL server database session contextual information. The following is the list of session context types:

- Hosts
- Program Names
- NT Domains
- NT Users
- Login Names
- Original Login Names
- Login Databases

6. After selecting the context type and an instance of it, click **Details**.

A dialog is displayed showing session/object details on one tab and device details on another tab.

## Viewing database storage details

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Database**.

3. Select an environment, click , and then click **VIEW ALL DETAILS**.


A screen is displayed with five tabs. The contents of the **Performance** tab are displayed by default.

4. Click the **Storage** tab.

The **Storage** tab consist of three panels: a **Database Properties** panel displaying Instance, version and last mapping date information; a **Database Array Properties** panel displaying the allocated capacity of the array; and a **Devices** panel listing the devices and their associated storage group, service level, capacity (in GB) and host device.

## Viewing database details

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Database**.
3. Select an environment, click , and then click **VIEW ALL DETAILS**.

A screen is displayed with five tabs. The contents of the **Performance** tab are displayed by default.

4. Click the drop-down menu to select the time range .
5. Click the **Databases** tab.

The **Databases** tab consists of a view displaying a list of environment names and their associated properties: DB RT Status, DB Read RT(ms), DB IOPs and DB MB/s.

## Adding hints


### Before you begin

- The storage system must be a hybrid array running HYPERMAX OS 5977 or higher.
- The database must be an Oracle or MS SQL Server database.
- DSA is accessible by users as described in [Roles](#) on page 70.

A hint helps to assure the best response time for mission critical processes or database objects.

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Databases** to view the full list of monitored environments and their associated properties.
3. Do one of the following:
  - Double click on a database name instance to navigate to the view all details section. The performance tab is displayed.

- Click on a database name instance and click  to open the details panel. Click **VIEW ALL DETAILS** to go to the **Performance** section for the monitored environment.
4. Click the **Analytics** tab.
  5. Select one, more than one, or all database objects and click **Add to Hint** to open the Hint wizard.
  6. Select **New** or **Existing Hint Type** and click **Next**.
  7. Type the hint name.
  8. Select the **Hint Priority**. Possible values are 1 through 3, with 1 being the highest.
    - **Priority 1**—Simulates a diamond service level. It marks all extents as active including the inactive ones and move data to EFD ahead of time. The primary use case for this priority is a periodic process such as end of month process where the assumption is that most of the data is inactive during the month.
    - **Priority 2**—Simulates a Platinum (OLTP) service level. It marks all extents as active (even the inactive ones) but there is no immediate promotion. The primary use case is ensuring that a given process receives better response time than other processes in the storage group even if the data was inactive.
    - **Priority 3**—Simulates a Platinum (DSS) service level change response time target without marking the inactive extents. The primary use case is to ensure that a given process receives better response time than other processes in the storage, yet it is less powerful than priority 2 as it only promotes active data.
  9. Select the **Hint Type**. Possible values are:
    - **One Time**—Type a value for the start date and time and the end date and time.
    - **Ongoing**—Hints are applied for the selected objects on an ongoing 24/7 basis.
    - **Recurrence** — Select the recurrence pattern.
      - **Daily**—Occurs on a daily basis.
      - **Weekly**—Select one or more days of the week.
      - **Monthly**—Type the day of the month.
  10. Type the values for the recurrence range—the start date and time and the end date and time.
  11. Optional: Deselect database objects to add to the hint.  
Note: there has to be at least one database object associated with the hint.
  12. Click **FINISH**.

## Viewing hints

### Before you begin

To view hint logs, refer to [Viewing hint logs](#) on page 858.

- The storage system must be a hybrid array running HYPERMAX OS 5977 or higher.

- The database must be an Oracle or MS SQL Server database.
- DSA is accessible by users as described in [Roles](#) on page 70.

A hint helps to assure the best response time for mission critical processes or database objects.

For an introduction to Database Storage Analyzer, see [Introduction](#) on page 834

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Hinting** to view the hints relating to all environments.  
You can also view hints for a specific environment by navigating to the view all details section and clicking on the **Hinting** tab.
3. Click **Hinting** to view all the hints relating to all the environments and their associated properties.

The following properties display for each hint:

- **Hint Name**—Name of the hint.
- **Environment Name**—Name of the database associated with the hint.
- **Host**—Name of the database host associated with the hint.
- **Priority**—Priority associated with the hint.
- **Status**—Status of the hint when it was last run.
- **Symmetrix ID**—Identity of the current Symmetrix that the database is running on.

The following controls are available:

- **Modify**—[Editing hints](#) on page 856
- **Enable**—[Enabling hints](#) on page 857
- **Disable**—[Disabling hints](#) on page 857
- **Remove**—[Removing hints](#) on page 858

## Editing hints

### Before you begin

- The storage system must be a hybrid array running HYPERMAX OS 5977 or higher.
- The database must be an Oracle or MS SQL Server database.
- DSA is accessible by users as described in [Roles](#) on page 70.

A hint helps to assure the best response time for mission critical processes or database objects.

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Hinting** to view the hints relating to all environments.
3. Click the hint that you wish to edit and click **Modify** to open the Hint wizard.
4. Optional: Modify the hint name.



5. Optional: Modify the **Hint Priority**. Possible values are 1 through 3, with 1 being the highest.
6. Optional: Modify the **Hint Type**. Possible values are:
  - **One Time** — Type a value for the start date and time and the end date and time.
  - **Ongoing** — Hints are applied for the selected objects on an ongoing 24/7 basis.
  - **Recurrence** — Select the recurrence pattern.
    - **Daily**—Occurs on a daily basis.
    - **Weekly**—Select one or more days of the week.
    - **Monthly**—Type the day of the month.
7. Optional: Modify the values for the recurrence range—the start date and end date.
8. Optional: Deselect database objects from the hint.  
Note: there has to be at least one database object associated with the hint.
9. Click **FINISH**.

## Enabling hints

### Before you begin

- The storage system must be a hybrid array running HYPERMAX OS 5977 or higher.
- The database must be an Oracle or MS SQL Server database.
- DSA is accessible by users as described in [Roles](#) on page 70.
- The hint is disabled.

A hint helps to assure the best response time for mission critical processes or database objects.

### Procedure

1. Click **HOME**.
2. Select **DATABASES > Hinting** to view the hints relating to all environments.
3. Select the hint that you wish to enable and click **Enable**.
4. Click **OK** to confirm the operation.

## Disabling hints

### Before you begin

- The storage system must be a hybrid array running HYPERMAX OS 5977 or higher.
- The database must be an Oracle or MS SQL Server database.
- DSA is accessible by users as described in [Roles](#) on page 70.
- The hint is enabled.

A hint helps to assure the best response time for mission critical processes or database objects.

**Procedure**

1. Click **HOME**.
2. Select **DATABASES > Hinting** to view the hints relating to all environments.
3. Select the hint that you wish to enable and click **Disable**.
4. Click **OK** to confirm the operation.

## Removing hints

**Before you begin**

- The storage system must be a hybrid array running HYPERMAX OS 5977 or higher.
- The database must be an Oracle or MS SQL Server database.
- DSA is accessible by users as described in [Roles](#) on page 70.

A hint helps to assure the best response time for mission critical processes or database objects.

**Procedure**

1. Click **HOME**.
2. Select **DATABASES > Hinting** to view the hints relating to all environments.
3. Select the hint that you wish to remove and click **Remove**.
4. Click **OK** to confirm the operation.

## Viewing hint logs

**Before you begin**

- The storage system must be running HYPERMAX OS 5977 or higher.
- The database must be an Oracle or MS SQL Server database.

A hint helps to assure the best response time for mission critical processes or database objects.

**Procedure**

1. Click **HOME**.
2. Select **DATABASES > Hinting** to view the hints relating to all environments.  
  
You can also view hint logs for a specific environment by navigating to the view all details section and clicking on the **Hinting** tab.
3. Select **Logs** from the drop down menu to view the full list of view the full list of hint logs and their associated properties.

You can also view hints for a specific environment by navigating to the view all details section and clicking on the Hinting tab.

The following properties display for each hint log:

- **Sent Time**—Time that the hint was sent to the storage system.
- **Hint Name**—Name of the hint.
- **Environment Name**—Name of the database associated with the hint.

- **Host**—Name of the database host associated with the hint.
- **Priority**—Priority associated with the hint.
- **Symmetrix ID**—Identity of the current Symmetrix that the database is running on.
- **Start Date**—Start date and time.
- **End Date**—End date and time.
- **Status**—Status of the hint when it was last run.
- **Description**—This lists the name of the user that created the hint.
- **Size(GB)**—The size of the log.
- **Hint ID**— Hint identity received from the storage system.

## Hint operations

- Viewing hints (see [Viewing hints](#) on page 855).
- Viewing hint logs (see [Viewing hint logs](#) on page 858).



# CHAPTER 9

## VMware

Topics include:

• <a href="#">Understanding Unisphere support for VMware</a> .....	862
• <a href="#">Viewing vCenters and ESXi information</a> .....	862
• <a href="#">Registering vCenter/ESXi</a> .....	863
• <a href="#">Editing vCenter/ESXi</a> .....	864
• <a href="#">Unregistering vCenter/ESXi</a> .....	864
• <a href="#">Rediscover vCenter/ESXi</a> .....	865
• <a href="#">Viewing ESXi server details</a> .....	865
• <a href="#">Viewing ESXi server masking views</a> .....	866
• <a href="#">Viewing ESXi server performance details</a> .....	866
• <a href="#">Viewing ESXi server virtual machines details</a> .....	867
• <a href="#">Viewing ESXi server virtual machine disks</a> .....	868

## Understanding Unisphere support for VMware

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

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

Once VMware information is added by a user, all other users of Unisphere are able to access this information.

The minimum version number supported by vCenter is version 5.5.

## Viewing vCenters and ESXi information

For an overview of Unisphere support for VMware, see [Understanding Unisphere support for VMware](#) on page 862.

This view displays a list view containing all the ESXi Hosts that will have relationships with a local storage system discovered by Unisphere.

A relationship between a storage system and a ESXi Host is defined as the storage system having a masking view configured with a storage group that is visible in ESXi Host Storage Adapters (Port Group configuration).

### Procedure

1. Select **VMWARE > vCenters and ESXi** from the menu bar on the left side of the screen.

The vCenters and ESXi list view is displayed. It displays the following properties:

- **ESXi**- ESXi name.
- **vCenter**- IP address of the vCenter.
- **SG Compliance**- Displays the value of Storage Group (SG) compliance of the worst performing SG that's associated with the ESXi Server.
- **Front End**- Displays the value of the worst performing Front End (FE) Port Busy for all arrays associated with the ESXi Server.
- **Discovered Time**- Date and time of discovery.

2. Click .

The vCenters and ESXi details view is displayed. It displays the following properties:

- **Name**- ESXi name.
- **Version**- ESXi server version.
- **Build Number**- ESXi server build number.
- **CPU Cores**- ESXi server hardware CPUs configuration.
- **Service Console Memory**- ESXi server hardware service console memory configuration.

- **Memory**- ESXi server hardware total memory configuration.

The following controls are available:

- **Register vCenter/EXSI Server** - [Registering vCenter/ESXi](#) on page 863
- **Edit** - [Editing vCenter/ESXi](#) on page 864
- **Rediscover** - [Rediscover vCenter/ESXi](#) on page 865
- **Unregister vCenter/EXSI Server** - [Unregistering vCenter/ESXi](#) on page 864

3. Optional: Choose a day or time to retrieve performance related data to the FE Ports over the selected period.

Example: ESXi Server A has five storage groups that are associated with two arrays (Array 1 and Array 2). The 5 SGs have the following compliance status:

- Three are stable
- One is marginal
- One is underperforming

Array 1 and Array 2 have the following Front End (FE Port) status: Two are stable.

In the above scenario, the following is displayed for ESXi Server A:

- SG Compliance: Underperforming (Red icon)
- FE Port : Stable (Green icon)

## Registering vCenter/ESXi

For an overview of Unisphere support for VMware, see [Understanding Unisphere support for VMware](#) on page 862.

You can register a new vCenter/ESXi Server by selecting "Register vCenter/ESXi Server" action in the VMWare view or by clicking the link for "Register vCenter/ESXi Server" in the Get Started section (if no servers are added yet).

---

### Note

This operation may take a considerable length of time for a large vCenter.

---

To register a vCenter/ESXi Server instance:

### Procedure

1. Select **VMWARE > VCenters and ESXi** from the menu bar on the left side of the screen.  
  
The vCenters and ESXi list view is displayed or the Getting Started view is displayed.
2. Click **Register vCenter/EXSI Server**.
3. Type the fully qualified name or IP address of the server (Server/IP Address).
4. Type the username used to connect to the server.
5. Type the password used to connect to the server.
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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

---

#### Note

This operation may take a considerable length of time for a large vCenter.

---

## Editing vCenter/ESXi

To edit an existing vCenter/ESXi Server instance:

### Procedure

1. Select **VMWARE > VCenters and ESXi** from the menu bar on the left side of the screen.

The vCenters and ESXi list view is displayed.

2. Select a vCenter/ESXi Server and click **Edit**.
3. Edit the fully qualified name or IP address of the server (Server/IP Address).

In case the user edits the server name to a totally new server, Unisphere removes the old server details and adds the new server details.

4. Edit the username used to connect to the server.

The user must have root privileges on the server.

5. Edit the password used to connect to the server.

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 920 and [Previewing jobs](#) on page 920.
- Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Unregistering vCenter/ESXi

You can unregister an individually added ESXi but if you attempt to remove an ESXi that was discovered under a vCenter, then that whole vCenter and all associated ESXis will be removed.


To unregister an existing vCenter/ESXi Server instance:

### Procedure

1. Select **VMWARE > VCenters and ESXi** from the menu bar on the left side of the screen.

The vCenters and ESXi list view is displayed.

- 2.

Select a vCenter/ESXi Server instance, click  , and click **Unregister vCenter/EXSI Server**.

A warning is displayed indicating that this operation will remove the ESXi or all ESXis under the selected vCenter.



3. Click **OK**.

## Rediscover vCenter/ESXi


### Procedure

1. Select **VMWARE > VCenters and ESXi** from the menu bar on the left side of the screen.  
The vCenters and ESXi list view is displayed.
2. Click **Rediscover**  
A dialog is displayed asking you to confirm that you want to rediscover all vCenters/ESXi servers.
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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to perform the operation now.

## Viewing ESXi server details

This tabbed view contains properties of the ESXi Server and also contains information on important storage performance key performance indicators (KPIs).

### Procedure

1. Select **VMWARE > VCenters and ESXi** from the menu bar on the left side of the screen.
2. Select a vCenter/ESXi Server instance and click .
3. Click **VIEW ALL DETAILS**.  
A view with multiple tabs is displayed.  
For information on the **Details** tab, see below.  
For information on the **Masking View** tab, see [Viewing ESXi server masking views](#) on page 866.  
For information on the **Virtual Machines** tab, see [Viewing ESXi server virtual machines details](#) on page 867.  
For information on the **Performance** tab, see [Viewing ESXi server performance details](#) on page 866.
4. Click the **Details** tab.  
The following properties are displayed:
  - **Name**- ESXi name.
  - **Discovered time**- Date and time of discovery.
  - **Build Number**- ESXi server build number.
  - **Version**- ESXi server version.
  - **CPU Count**- ESXi server hardware CPU count.

- **Service Console Memory**- ESXi server hardware service console memory configuration.
- **Total Memory**- ESXi server hardware total memory configuration.
- **Virtual Machines** - Total number of virtual machines belonging to the ESXi Server


The following array related storage properties are displayed:

- **Masking Views** - The total number of masking views found for the ESXi server.
- **Storage Groups** - The total number of storage groups found for the ESXi server.
- **Capacity (GB)** - The capacity for all the Storage Groups found for the ESXi Server (this is calculated by doing a sum of all the SGs)
- **Allocated Capacity** - The allocated capacity for all the Storage Groups found for the ESXi Server (this is calculated by doing a sum of all the SGs)

## Viewing ESXi server masking views

This tabbed view contains all the masking views associated with the ESXi server.

### Procedure

1. Select **VMWARE > VCenters and ESXi** from the menu bar on the left side of the screen.
2. Select a vCenter/ESXi Server and click .
3. Click **VIEW DETAILS**.

A view with multiple tabs is displayed.

4. Click the **Masking View** tab.

The following properties are displayed:

- **Masking View** - Name of the masking view.
- **vCenter** - The name or IP address of the vCenter associated with ESXi server.
- **Array** - The Array associated with the masking view.

The following controls are available:


- **Rediscover** - [Rediscover vCenter/ESXi](#) on page 865
- **View Path Details** - [Viewing masking view connections](#) on page 309

## Viewing ESXi server performance details

This tabbed view contains the performance details for all the Storage Groups related with the ESXi Server.

There are two sections in this tab. The top section comprises a list view that displays all the Storage Groups found and their performance metrics. On the bottom section there is the noisy neighbor information that is displayed when you select a Storage Group from the top section.

### Procedure

1. Select **VMWARE > VCenters and ESXi** from the menu bar on the left side of the screen.
2. Select a vCenter/ESXi Server and click .
3. Click **VIEW ALL DETAILS**.

A view with multiple tabs is displayed.

4. Click the **Performance** tab.

The top section is a list view containing the Storage Groups. The following properties are displayed:

- **Storage Group** - Name of the storage group. Clicking on the name will take you to the Storage Group Performance dashboard.
- **Array ID** - Array associated with the storage group.
- **Host I/Os per sec** - Storage Group Host I/Os/sec
- **Host MBs per sec** - Storage Group Host MBs/sec
- **Read RT (ms)** - Storage Group Read RT time
- **Write RT (ms)** - Storage Group Write RT time
- **SG Compliance** - Displays the value of Storage Group (SG) compliance.

5. Select a storage group by clicking on the checkbox.

The Noisy Neighbor section displayed. It is composed of the following performance data for the selected storage group:

- **FE Directors** details - Name, % busy, and queue depth utilization.
- **FE Port** details - Name, % busy, and host I/Os per second.
- **Related SGs** - Name, response time, host I/Os per second, and host MBs per second.


The following control is available:

- **Rediscover** - [Rediscover vCenter/ESXi](#) on page 865

## Viewing ESXi server virtual machines details

This tabbed view contains details about the virtual machines (VMs) under the ESXi Server.

### Procedure

1. Select **VMWARE > VCenters and ESXi** from the menu bar on the left side of the screen.
2. Select a vCenter/ESXi Server instance and click .
3. Click **VIEW ALL DETAILS**.

A view with multiple tabs is displayed.

4. Click the **Virtual Machines** tab.

The following properties are displayed:


- **Name**

- **Powered**
- **CPU Count**
- **Total Memory (MB)**
- **OS Name**

The following control is available:

- **Rediscover** - [Rediscover vCenter/ESXi](#) on page 865

5.

Click  to view the following additional VM properties:

- **Address** - IP address.
- **Host Name** - Host name.
- **Virtual Disks** - A number indicating the number of virtual disks of this virtual machine. You can click the link in order to navigate to a list showing all the disks associated with the virtual machine.

## Viewing ESXi server virtual machine disks

This tabbed view contains details about the virtual machine (VMs) disks under the ESXi Server.

### Procedure

1. Select **VMWARE > VCenters and ESXi** from the menu bar on the left side of the screen.


2. Select a vCenter/ESXi Server instance and click .


The vCenters and ESXi details view is displayed. It displays the following properties:

3. Click **VIEW ALL DETAILS**.

A view with multiple tabs is displayed.

4. Click the **Virtual Machines** tab.

5. Select a virtual machine, click , and click the number next to **Virtual Disks**.

6. Select a virtual disk and click  to view the following additional physical disk properties.

- **Capacity** - Physical Disk capacity.
- **Model** - Physical disk model.
- **WWN** - Physical disk WWN.
- **Name** - Physical disk name.
- **Vendor** - Physical disk vendor name.
- **Path** - Physical Disk Path.
- **Datastore** - Datastore for physical disk.

- **Type** - Physical Disk type.



# CHAPTER 10

## System Management

• Viewing Storage System details.....	872
• Setting system attributes.....	874
• Using the Emulation Management wizard.....	877
• Setting CPU I/O resource distribution.....	877
• Setting logging level preferences .....	878
• Understanding eNAS.....	878
• Viewing the system audit log.....	889
• Viewing system hardware.....	891
• Managing jobs.....	918
• Understanding licenses.....	925
• Understanding access controls.....	931
• Understanding dynamic cache partitioning.....	943
• System management - iSCSI.....	951


## Viewing Storage System details

This procedure explains how to view storage system details. You can also view this information by clicking **Dashboard > System Health > Actions > VIEW SYMMETRIX PROPERTIES**.

### Procedure

1. Select the storage system.
2. Select **System > Symmetrix Properties**.

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 . 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** —The product model.
- **HyperMax OS or Enginuity or PowerMaxOS**—The microcode version number.
- **Microcode Date/Patch Date** —The microcode date or patch date.
- **Capacity**—Capacity of the storage system.
- **Used Capacity**—not applicable to storage systems running HYPERMAX OS 5977 or greater.
- **Free capacity**— not applicable to storage systems running HYPERMAX OS 5977 or greater.
- **Actual Capacity**—not applicable to storage systems running HYPERMAX OS 5977 or greater.
- **Cache Size (Mirrored)**—Cache size.
- **Number of Front End Directors**—Number of Front End Directors.
- **Number of Back End Directors**—Number of Back End Directors.
- **Number of RDF Directors**—Number of RDF Directors.
- **Max Number of Hypers per Disk**—Maximum number of hypers per disk.
- **Device Masking (VCM) Config State**—Device Masking VCM configuration state.
- **VCMDB Access Restricted State**—VCMDB Accesss Restricted State.
- **Device Masking (ACLX) Config State**—Device Masking ACLX configuration state.
- **ACLX Volume LUN Address**—ACLX Volume LUN Address.
- **Access Control Configuration State**—Access Control Configuration state.
- **Configuration Change State**—Configuration Change state.
- **Disk Group Assignments**—Disk Group Assignments.
- **Parity RAID Configuration**—Parity RAID configuration.
- **RAID Configuration**—RAID configuration.



- **Pav Mode**—PAV Mode.
- **Pav Alias Limit**—PAV Alias Limit.
- **SDDF Configuration State**—SDDF Configuration state.
- **WORM Configuration Level**—WORM configuration level.
- **WORM Characteristics**—WORM characteristics.
- **Switched RDF Configuration State**—Switched RDF configuration state.
- **Concurrent RDF Configuration State**—Concurrent RDF configuration state.
- **Dynamic RDF configuration State**—Dynamic RDF configuration state.
- **Concurrent Dynamic RDF State**—Concurrent Dynamic RDF state.
- **RDF Data Mobility Configuration State**—RDF Data Mobility Configuration State .
- **SRDF/A Maximum Host Throttle (Secs)**—SRDF/A Maximum Host Throttle (seconds).
- **SRDF/A Maximum Cache Usage (Percent)**—SRDF/A Maximum Cache Usage (Percent).
- **Multi LRU Device Assignment**—Multi LRU Device Assignment.
- **Number of Available Cache Slots**—# of Available Cache Slots.
- **Max Number of DA Write Pending Slots**—Max # of DA Write Pending Slots.
- **Max Number of System Write Pending Slots**—Max # of System Write Pending slots.
- **Max Number of Device Write Pending Slots**—Max # of Device Write Pending slots.
- **Max Number of Replication Slots**—not applicable to storage systems running HYPERMAX OS 5977 or greater.
- **Symmetrix Last IPL Time (Cold)**—Symmetrix Last IPL Time (Cold).
- **Symmetrix Last Fast IPT Time (Hot)**—Symmetrix Last Fast IPT Time (Hot).
- **Symmetrix Alerts (Enabled/Disabled)**—Symmetrix Alerts (Enabled/Disabled).
- **Max Capacity of DSE reservation**—not applicable to storage systems running HYPERMAX OS 5977 or greater.
- **Symmetrix Priority Controls**—not applicable to storage systems running HYPERMAX OS 5977 or greater.
- **Hot Swap Policy**—Hot Swap Policy.
- **Symmetrix Disk Library**— Symmetrix Disk Library.
- **FBA Geometry Emulation**— FBA Geometry Emulation.
- **3 Dynamic Mirrors**—3 Dynamic Mirrors.
- **Cache Partitioning**—Cache Partitioning.
- **Auto Meta**—Auto Meta.
- **Minimum Auto Meta Size (Cyl)**—Minimum Auto Meta Size (Cyl).

- **Auto Meta Member Size (Cyl)**—not applicable to storage systems running HYPERMAX OS 5977 or greater.
- **Auto Meta Configuration**—Auto Meta Configuration.
- **Disk Service State**—Disk Service State.
- **Front door LED status**—Front door LED status.
- **Symmetrix Data Encryption**—Status of Symmetrix Data Encryption (Enabled/Disabled).
- **Replication Cache Usage (Percent)**—not applicable to storage systems running HYPERMAX OS 5977 or greater.


## Setting system attributes

### Before you begin

Depending the storage system operation environment, some of the following attributes may not apply.

To set system attributes:

### Procedure

1. Select  to open the **Settings** panel.
2. Select **Management > Symmetrix Attributes** to open the **Symmetrix Attributes** page.  
  
The Symmetrix attributes information is displayed for the first of the listed storage arrays.
3. Select the ID of the storage system to be updated from the drop down menu. Depending on the storage system environment, some might not apply.
4. 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.

- Enable Auto Meta:  
This panel only displays for storage systems running Enginuity 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 5876 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.

### 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 storage system.

##### StandardPAV

Volumes with static aliasing.

##### DynamicStandardPAV

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

Indicates 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 Percent**

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 946.

**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.

**Maximum number of Partitions**

Lists the maximum number of partitions.

**Minimum Allowed Target %**

Lists the minimum allowed target percentage.

**Maximum Allowed Target %**

Lists the maximum allowed target percentage.

**XRC Partition State**

Lists the state of the XRC partition.

**XRC Partition Name**

Displays the XRC partition name.

**Time of Last Modification**

Displays the time that the last modification occurred.

- PowerPath

**Host Registration**

Checkbox used to select host registration for PowerPath.

**Initiator Registration**

Checkbox used to select initiator registration for PowerPath.

5. Click **APPLY**.

## Using the Emulation Management wizard

### Before you begin

- The storage system must be running HYPERMAX OS 5977 Q1 2016, or higher.
- StorageAdmin or Admin authorization rights are required.
- A free director must be available at the required slot.
- The addition of directors of following types is not supported:
  - IM—Infrastructure Management
  - ED—Enginuity Data Services
  - DS—SAS back-end
  - DA—Fibre back-end
  - DX—External storage back-end
  - EF—FICON front-end

You can use the Emulation Management wizard to add and remove director emulations.

### Procedure

1. Select the storage system.
2. In the **Dashboard**, click the **System Health** tab.
3. In the **Actions** panel, click **Manage Emulation**.
4. On the **Options** page, select one of the following:
  - Add Director Emulation
    - a. Click **Next**.
    - b. On the **Configure** page, select values for the following:
      - **Director Slot**
      - **Director Emulation**
    - c. Click **Next**.
  - Remove Director Emulation
    - a. Click **Next**.
    - b. On the **Configure** page, select the emulation to remove.
    - c. Click **Next**.
5. 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 920 and [Previewing jobs](#) on page 920.

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

## 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 [Setting system attributes](#) on page 874.

**Procedure**

1. Select the storage system.
2. Select **System > Hardware > RDF Directors**.
3. Select an RDF director and click **CPU I/O Setting**.
4. Type values for any number of the following parameters. Note that the percentages must add up to 100.
  - **Synchronous I/O Percent**—Type the percentage of director resources to allocate for synchronous I/Os.
  - **Asynchronous I/O Percent**—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.
5. To reset the parameters to their defaults, click **Reset**.
6. Click **OK**.

## Setting logging level preferences

This procedure explains how to configure the level of message to maintain in the Unisphere log.

**Procedure**

1. Select **Support** and in the **Actions** panel, click **Modify Server Logging**.
2. Select a **Logging Level**.
3. Click **OK**.

## Understanding eNAS

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

With this integrated storage solution, Unisphere StorageAdmin provision 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 provides the following features to support eNAS:

**File System dashboard**

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

**Provision Storage for File wizard**

Allows you to provision storage to eNAS data movers.

**Launch Unisphere for VNX**

Allows you to link and launch Unisphere for VNX.

## Discovering eNAS control stations

### Before you begin

- To perform this operation, you must be an Administrator.
- The storage system must be running HYPERMAX OS 5977 or higher.

This procedure explains how to connect to an eNAS control station for the first time.

### Procedure

1. Select the storage system.
2. Select **System > File** to open the **File Dashboard**.
3. If eNAS is not authenticated, click the **Open Authentication dialog** link.
4. Type the username and password you use to access the eNAS control station and click **OK**.

## Managing 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.

The File dashboard provides you with a single place from which to monitor and manage integrated VNX file services

### Procedure

1. To access the File dashboard:
  1. Select the storage system.
  2. Select **System > File**.
  3. If prompted, type the username and password you use to access the eNAS control station, and click **OK**.

The File 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.

For more information about the elements of the File dashboard, refer to [The File Dashboard](#).

## The File dashboard

The File dashboard is organized into the following panels:

- Summary
- Actions
- Capacity
- Control Stations
- Data Movers
- File Storage Alerts

### 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.
- **File Storage Pools**—File storage pools associated with the file storage groups.

### Actions

Links to common file storage tasks, including:

- [Provisioning storage for file](#) on page 880
- [Launching Unisphere for VNX](#) on page 882

### Capacity

Displays the following capacity information:

- **Usable Capacity**—Free versus total capacities of the file storage groups.
- **File Capacity**—Free versus total capacities for the file systems associated with the file storage groups on the storage system.

### Control Stations

Displays the name and status of the control stations.

### Data Movers

Displays the status of the control station's data movers:

- **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:

- **State**—State of the alert.
- **Severity**—Alert severity. Possible values are:
  - (1) Fatal
  - (2) Critical
  - (3) Warning
  - (4) Information
  - (5) Normal
- **Type**—State of the alert.
- **Description**—Description of the alert.
- **Created**—Date/time the alert was issued.

## Provisioning storage for file

### 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 provision storage to eNAS datamovers.

To provision storage for file:



## Procedure

1. Select the storage system.
2. Select **System > File** to open the **File Dashboard**.
3. If prompted, type the username and password you use to access the eNAS control station, and click **OK**.
4. In the **Actions** panel, click **Provision Storage for File** to open the Provision Storage for File wizard.
5. Supply a **Storage Group Name**.


Each storage group name is unique among the storage groups on the storage system and has up to 64 characters. Use only alphanumeric characters, underscores ( \_ ), and dashes ( - ) in the name. Storage group names are case-insensitive.


6. To select a **Storage Resource Pool** other than the default (DEFAULT\_SRP), click on the drop-down menu and select one. To create the storage group outside of FAST control, select **None**.
7. 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. Available values are:

Service level	Performance type	Use case
Diamond	Ultra high	HPC, latency sensitive
Platinum	Very high	Mission critical, 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.

For all-flash storage systems, the only service level available is Diamond and it is selected by default.

8. Refine the service level by selecting the **Workload Type** to assign to it.
9. Type the number of **Volumes** and select the **Capacity** of each.
10. (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**.  
You can set this attribute 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**.
  - To persist preallocated capacity through reclaim or copy, select that checkbox.

11. (Optional) To add a child storage group, do one of the following: click  and
  - On all-flash storage systems, click **Add Storage Group**.
  - On all other storage systems, click **Add Service Level**.

Specify a **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.
12. (Optional) To set host I/O limits for the storage groups, click **Set Host I/O Limits** in the **Summary** tab to open the **Host I/O Limits** dialog box. For information about setting the limits, refer to the dialog's help page.
13. Click **Next**.
14. (Optional) On the **Review** page of the wizard, follow these steps to manually set the LUN addresses for the masking operation:
  - a. Click **Set Dynamic LUNs** to open the **Set Dynamic LUN Addresses** dialog.
  - b. Note 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. Optional: Click **Enable compliance alerts**.
  - d. Click **OK**.
15. Click **Run Now** or **Add To Job List**.

## Launching 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.

### Procedure

1. Select the storage system.
2. Select **System > File**.
3. If prompted, type the username and password you use to access the eNAS control station, and click **OK**.
4. In the **Actions** panel, click **Launch Unisphere for VNX**.

## Managing 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.

This procedure explains how to manage file storage groups.

### Procedure

1. Select the storage system.

2. Select **System > File** to open the **File Dashboard**.
3. In the **Summary** panel, click **File Storage Groups**.
4. Select one or more storage groups and click one of the following, depending on the operation you want to perform:
  -  — [Viewing storage group details](#) on page 138
  - **Modify**—[Modifying storage groups](#) on page 119
  - **Provision Storage for File**—[Provisioning storage for file](#) on page 880
  - **FAST Array Advisor**
  - **Change SRP**—[Changing Storage Resource Pools for storage groups](#) on page 129
  - **Delete**—[Deleting storage groups](#) on page 131
  - **Start Allocate/Free/Reclaim**—[Managing thin pool allocations](#) on page 244
  - **Stop Allocated/Free Reclaim**—[Managing thin pool allocations](#) on page 244
  - **Convert to Cascaded**—[Converting storage groups to cascaded](#) on page 129
  - **Set Volume Status**—[Setting volume status](#) on page 435

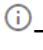
## Managing file masking views

### 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.

This procedure explains how to manage file masking views.

### Procedure

1. Select the storage system.
2. Select **System > File** to open the **File Dashboard**.
3. In the **Summary** panel, click **File Masking Views**.
4. Select one or more masking views and click one of the following, depending on the operation you want to perform:
  -  — [Viewing masking view details](#) on page 311
  - **View Path Details**—[Viewing masking view connections](#) on page 309
  - **Delete**—[Deleting masking views](#) on page 308

## Viewing 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.

This procedure explains how to view VNX file systems.

### Procedure

1. Select the storage system.
2. Select **System > File** to open the **File Dashboard**.
3. 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 Group(s)**—Storage group(s) 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:



— [Viewing file system details](#) on page 884

## Viewing file system details

### 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.

This procedure explains how to view details on VNX file systems.

### Procedure

1. Select the storage system.
2. Select **System > File** to open the **File Dashboard**.
3. In the **Summary** panel, click **File Systems** to open the File Systems list view.

4. Select a file system and click  to open its Details view.

For more information about the File Systems details view, refer to [File System Details](#) on page 884.

## File System Details

The File System Details view displays the following properties.

- **File System Name** — Name of the file system.
- **File System Unique ID** — Unique VNX ID for the file system.
- **File System Total Capacity (GB)** — Size of the file system.
- **File System Free Capacity (GB)** — Available space on the file system.
- **File System Block Size (Bytes)** — 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.

- **Storage Groups** — Storage groups on the file system.
- **Storage Pool** — Storage pools on the file system.

## Viewing file system storage pools

### 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 view storage pools associated with VNX file systems.

### Procedure

1. Select the storage system.
2. Select **System > File** to open the **File Dashboard**.
3. In the **Summary** panel, click **File Storage Pools**.

The following properties display:

- **Storage Pool Name**—Name of the storage pool.
- **Instance ID**—Long name description of the storage pool.
- **#File Systems**—File system using the storage pool in a VNX context.
- **Health State**—Health state of the storage pool.
- **Storage Group**—Name of the storage group.
- **Allocated %**—Percentage allocated for the storage pool.
- **Capacity (GB)**—Capacity of the storage pool in gigabytes.

The following controls are available:

-  — [Viewing file system storage pool details](#) on page 885
- **Expand SG for File**—[Modifying storage groups](#) on page 119


## Viewing file system storage pool details

### 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 view storage pools associated with VNX file systems.

### Procedure

1. Select the storage system.
2. Select **System > File** to open the **File Dashboard**.
3. In the **Summary** panel, click **File Storage Pools**.
4. Select a file storage pool and click .

### File system storage pool details

The following properties are displayed:

- **Storage Pool Name** — Name of the storage pool.
- **Instance ID** — Long name description of the storage pool.
- **Health State** — Current health of the storage pool. Values for this property are:
  - 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. Values for this property are:
  - Unknown
  - Not Available
  - Servicing
  - Starting
  - Stopping
  - Abandoned
  - Dormant
  - Completed
  - Migrating
  - Emigrating
  - Immigrating
  - Snapshotting
  - 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**
- **Number of File Systems** — Click on the link in this field to see the file systems in the storage pool. The following properties display:
  - **File System Name**—Name of the file system.
  - **Storage Group(s)**—Storage groups mapped to the file system.

- **Storage Pools**—Storage pools in the file system.
- **Allocated %**—Percentage of the file system allocated.
- **Capacity (GB)**—Total capacity in GB of the file system.

## Manage file storage alerts

### Viewing file storage alerts

#### 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.

This procedure explains how to view alerts related to file storage.

#### Procedure

1. Select the storage system.
2. Select **System > File** to open the **File Dashboard**.
3. The **File Storage Alerts** panel displays the following properties:
  - **State**—State of the alert. Available values are **New** or **Acknowledged**.
  - **Severity**—Alert severity. Values of this property are:
    - (1) Fatal
    - (2) Critical
    - (3) Warning
    - (4) Information
    - (5) Normal
  - **Type**—Type of alert. In this view, this property always contains File.
  - **Symmetrix**
  - **Object**—Object to which the alert is related.  
This field is blank for server alerts; server alerts are specific to the server or runtime environment and so 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.

To see all alerts, click **View All Alerts**.

### Viewing file storage alert details


#### 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.

This procedure explains how to view details on alerts related to file storage.

#### Procedure



1. Select the storage system.

2. Select **System > File** to open the **File Dashboard**.
3. In the **File Storage Alerts** panel, click **View All Alerts**.
4. Select an alert and click  to open the file storage Alert Details view.  
The following properties display:
  - **Alert ID**—Unique number assigned by Unisphere.
  - **State**—State of the alert. The value of this property is **New** or **Acknowledged**.
  - **Severity**—Alert severity. Values are:
    - (1) Fatal
    - (2) Critical
    - (3) Warning
    - (4) Information
    - (5) Normal
  - **Type**—Type of alert. The values of this property 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.

## Deleting file storage alerts

You can delete file storage alerts from the main **Alerts** list view or the **File Storage Alerts** view.

### Procedure

1. Do one of the following:
  - To delete a file storage alert from the main Alerts list view, complete the following steps:
    - a. Select the storage system.
    - b. Select **EVENTS > Alerts** to open the Alert list view.
    - c. Select one or more **File** type alerts and click .
  - To delete a file storage alert from the File Storage Alerts view, complete the following steps:
    - a. Select the storage system.
    - b. Select **System > FileeNAS** to open the **File Dashboard**.
    - c. In the **File Storage Alerts** panel, click **View All Alerts** to open the File Storage Alerts list view.
    - d. Select one or more alerts and click .



# Viewing the system audit log

## Before you begin

Because system audit logs can be very large, the list view opens pre-filtered to display the most recent 512 records.

The user must have Auditor or higher permissions.

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.

To view the system audit log:

## Procedure

1. Select the storage system.
2. Select **Events > Audit Log** to open the Audit Log list view.

The following properties are displayed:

### Record

Unique identifier for the audit entry.

### Date

Date the audit entry was made.

### Application

Application operating on the storage system.

### Action Code

Specific audit code for the operation 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.

### OS Type

Operating system running on the host.

### Activity ID

Activity ID for audit record.


To filter the log records, click  .

To view details of a specific audit log record, click  — [Viewing Symmetrix audit log details](#) on page 890 .

## Viewing Symmetrix audit log details

To view details on a specific audit log record:

### Procedure

1. Select the storage system.
2. Select **Events > Audit Log**.
3. Select a record and click  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

SYMAPI 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.

## Viewing system hardware

By selecting a storage system and clicking **System > Hardware**, you can view directors and available ports.

### Viewing available ports

#### Before you begin

---

##### Note

The **Available Ports** tab is not available on systems running Enginuity 5876.

---

#### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. Click the **Available Ports** tab to open the **Available Ports** list view.

This 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 GB/sec**

Transmission rate (input/output channel).

The following control is available:

- **Associate**—[Associating directors and ports](#) on page 912

## Viewing back-end directors

### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. Click the **BE Directors** tab 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.

#### % 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 5876.

#### Protocol

Transmission protocol for the director port.

#### Cores

Number of director cores. This field only displays for storage systems running HYPERMAX OS 5977 or higher.

#### Director Status

Director status.

#### Port Interface

Port interface ID. This field only displays for storage systems running Enginuity 5876.

#### Port


Port ID. This field only displays for storage systems running HYPERMAX OS 5977 or higher.

#### Port Status

Indicates whether the port is online or offline.


#### Alerts

Indicates if there are alerts associated with the director.

To see more information on a director's details and performance, select it and click .

## Viewing back-end director details

### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. Click the **BE Directors** tab to open the Back End Directors list view.
4. Select a director from the list and click  to open the Back End Director details view.

The following properties display under the **Details** tab:

#### 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.

The **Performance** tab links you to the performance monitor and analyze views for the director. This panel displays only when the Performance option is installed. Inactive links will be displayed if the selected storage system is not registered for data collection.

## Viewing external directors

### Before you begin

External directors are supported only on storage systems running HYPERMAX OS 5977 or higher.

### Procedure

1. Select the storage system.
2. In the **Dashboard**, click the **System Health** tab.
3. In the **Actions** panel, click **View Other Hardware** and then click the **External** tab to open the External Directors list view.

---

**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**


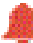



Number of available director cores.

**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 controls are available:

- **Associate**—[Associating directors with ports](#) on page 910 (This control only displays when the storage system is running HYPERMAX OS 5977 or higher.)
- **Disassociate**—[Disassociating directors and ports](#) on page 913 (This control only displays when the storage system is running HYPERMAX OS 5977 or higher.)

## Viewing external director details

### Procedure

1. Select the storage system.
2. In the **Dashboard**, click the **System Health** tab.
3. In the **Actions** panel, click **Other Hardware** and then click the **External** tab to open the External Directors list view.

4. Select the director and click  to open its details view.

The external director details view allows you to view details on an external director.

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.

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

This panel only displays when the Performance option is installed. Inactive links will be displayed if the selected storage system is not registered for data collection.

## Viewing system front-end directors

### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. Click the **FE Directors** tab 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.

**% Busy**

Percent of time the director is busy. This field only displays when the storage system is registered with the Performance component.

**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.

**ACLX**

Indicates whether access control logic is enabled or disabled.

**Cores**

Number of cores. This field appears only 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).

To see more information on a director's details and performance, select it and

click 

The following controls are available:

- **Enable**—[Enabling and disabling director ports](#) on page 913
- **Disable**—[Enabling and disabling director ports](#) on page 913
- **Set Port Attributes**—[Setting director port attributes](#) on page 910
- **Convert FA to RF**—[Converting directors](#) on page 909 (not applicable to storage systems running HYPERMAX OS 5977 or higher.)
- **Set ORS Ceiling**—[Setting Open Replicator ceiling](#) on page 494



- **Map**—[Mapping volumes](#) on page 192 (Only enabled when ACLX is disabled)
- **Unmap**—[Unmapping volumes](#) on page 193 (Only enabled when ACLX is disabled)
- **Associate**—[Associating directors with ports](#) on page 910 (Only displays when the storage system is running HYPERMAX OS 5977 or higher.)
- **Disassociate**—[Disassociating directors and ports](#) on page 913 (Only displays when the storage system is running HYPER MAX OS 5977 or higher.)

## Viewing system front end director details

### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. Click the **FE Directors** tab to open the Front End Directors list view.
- 4.

Select a director and click  to open the Front End Director details view.

The Front End Director details view allows you to view and manage a front end director.

The following properties display under the **Details** tab:

#### Director

Director name

#### Port

Port ID

#### 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.

#### 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 Set Addressing is enabled or disabled.

**VNX Attached**

Indicates whether VNX Attached is enabled or disabled.

**Avoid Reset Broadcasting**

Indicates whether Avoid Reset Broadcasting feature is enabled or disabled.

**Negotiate Reset**

Indicates whether Negotiate Reset feature is enabled or disabled. This feature is used for AS/400 systems only (default is off).

**Enable Auto Negotiate**

Indicates whether Enable Auto Negotiate feature is enabled or disabled.

**Environ Set**

Indicates whether environmental error reporting feature is enabled or disabled.

**Disable Q Reset on UA**

Indicates whether the Disable Q Reset on UA (Unit Attention) is enabled or disabled.

**Soft Reset**

Indicates whether Soft Reset feature is enabled or disabled for a Bull/ GCOS-7 host.

**SCSI 3**

Indicates whether SCSI-3 protocol is enabled or disabled. When disabled, the SCSI 2 protocol is supported.

**SCSI Support1(OS2007)**

Indicates whether SCSI Support1 is enabled or disabled.

**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.

**OpenVMS**

Indicates whether OpenVMS is enabled or disabled.

**Open Replicator Maximum Ceiling (MB)**

Maximum recopy ceiling value.

**Open Replicator Ceiling (%)**

Recopy ceiling value.

**Open Replicator Actual Ceiling (MB)**

Recopy actual ceiling value.

**Negotiated Speed (GB/Second)**

Negotiated speed in GB/Second.

The **Performance** panel links you to the performance monitor and analyze views for the director. This panel displays only when the Performance option is installed. Inactive links will be displayed if the selected storage system is not registered for data collection.

## Viewing RDF directors

**Procedure**

1. Select the storage system.
2. Select **System > Hardware**.
3. Click the **RDF Directors** tab 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.

**% 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.

**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.

To see more information on a director's details and performance, select it and click 


The following controls are available:

- **Enable**—[Enabling and disabling director ports](#) on page 913
- **Disable**—[Enabling and disabling director ports](#) on page 913

- **Convert RF to FA**—[Converting directors](#) on page 909 (not applicable to storage systems running HYPERMAX OS 5977 or higher)
- **CPU I/O Setting**—[Setting CPU I/O resource distribution](#) on page 877
- **Associate**—[Associating directors and ports](#) on page 912 (only applicable to storage systems running HYPERMAX OS 5977 or higher)
- **Disassociate**—[Disassociating directors and ports](#) on page 913 (only applicable to storage systems running HYPERMAX OS 5977 or higher)

## Viewing RDF director details

### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. Click the **RDF Directors** tab to open the RDF Directors list view.
4. Select a director from the list and click  to open its details view.

The RDF Director details view allows you to view RDF director details and performance.

The following properties display in the **Details** panel:

#### Director Identifier

Director name.

#### Type

Director type.

#### Director Status

Director status.

#### Port

Port ID.

#### WWN

World wide name.

#### RDF Groups

#### RDF (RA) Group Attributes - Farpoint

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.

**iSCSI 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.

The **Performance** panel links you to the performance monitor and analyze views for the RDF director. This panel displays only when the Performance option is installed. Inactive links will be displayed if the selected storage system is not registered for data collection.

## Viewing RDF director SRDF groups

### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. Click the **RDF Directors** tab to open the RDF Directors list view.
- 4.

Select a director from the list and click  to open its details view.

The RDF Director details view allows you to view RDF director details and performance.

The following properties display in the **Details** panel:

#### Director Identifier

Director name.

#### Type

Director type.

#### Director Status

Director status.

#### Port

Port ID.

#### WWN

World wide name.

#### RDF Groups

#### RDF (RA) Group Attributes - Farpoint

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.

**iSCSI 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.

The **Performance** panel links you to the performance monitor and analyze views for the RDF director. This panel displays only when the Performance option is installed. Inactive links will be displayed if the selected storage system is not registered for data collection.

5. Click on the number in the **RDF Groups** field to open the RDF Groups list view.

Use this view to view the SRDF groups mapped to the front-end director.

The following properties display:

**SRDF Group**

RDF group number.

**SRDF Group Label**

RDF group label.



**Remote SRDF group**

Remote RDF group number.

**Online****Volume Count**

Number of volumes in the group.

**SRDF Mode****Type**

Whether volumes in the group are dynamic or static.

**Transmit Idle**

Time the transmit cycle has been idle.

## Viewing IM directors

**Before you begin**

Infrastructure Manager (IM) directors are only supported on storage systems running HYPERMAX OS 5977 or higher.

**Procedure**

1. Select the storage system.
2. In the **Dashboard**, click the **System Health** tab.
3. In the **Actions** panel, click **View Other Hardware** and then click the **IM Directors** tab.

The following properties display in the IM Directors list view:

**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.
- **% Busy**—
- **Cores**—
- **Alerts** — Director alert status, as indicated by icon color:



One or more fatal alerts



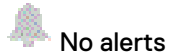
One or more critical alerts, with none higher



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 storage systems running Enginuity version 5977 or higher.

### Procedure

1. Select the storage system.
2. In the **Dashboard**, click the **System Health** tab.
3. In the **Actions** panel, click **Other Hardware** and then click the **EDS** tab to open the EDS 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.

#### % 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.

## Viewing failed drives

### Procedure

1. Select the storage system.
2. Select the **System Health** tab.
3. In the **Actions** panel, click **Run Drive Replacement**.

Use this view to view and replace failed disk drives.

The following properties display:

#### Disk Technology

Type of disk.

#### Capacity

Usable disk capacity.

#### Speed (RPM)

Physical disk revolutions per minute.

#### Failed Spare Drive Counts

Number of available spares.

#### Block Size


Size of each disk block.

The following control is available:

- **Replace**—[Replacing failed drives](#) on page 914

## Viewing mapped front-end volumes

### Procedure

1. Select the storage system.
2. Select **System > Hardware > FE Directors**.
3. Select a director and click .
4. In the **Details** tab, click on the number next to the **Volumes** field.

The following properties display:

#### Name

Assigned volume name.

#### Type

Type of volume.

#### Allocated %

**Capacity (GB)**

Volume capacity in Gigabytes.

**Status**

Volume status.

**Emulation**

Emulation type for the volume.

**SRDF Group**

**Host Paths**

To see more details of the volume, click



## Enginuity Warning dialog box

This dialog box displays when one or more storage systems on which you have admin rights is running an Enginuity version below the minimum suggested version. In which case, you should contact your service provider to schedule an upgrade.

In addition to this dialog box, an alert will also display for each storage system not running the minimum suggested Enginuity version. For more information on viewing alerts, refer to [Viewing alerts](#) on page 52.

This dialog displays the following properties:

**Symmetrix**

Storage system ID.

**Enginuity**

Enginuity version currently running on the storage system.

**Date**

Date the current Enginuity version was released.

**Target Enginuity**

Minimum suggested Enginuity version for the storage system.\*

**Hide Warning**

Select this option to no longer view this warning for the corresponding storage system.

\* This field will appear empty when the:

- Enginuity version currently running on the storage system is the same as the **Target Enginuity** version and it is more than **180 Days Old**.
- Enginuity version currently running on the storage system is lower than 5876 or higher than HYPERMAX OS 5977.
- Enginuity version currently running on the storage system is at the end of life (EOL). For more information on the EOL Enginuity versions, refer to the release notes.

\*\* This field will appear empty when the Enginuity version currently running on the storage system is lower than the **Target Enginuity**.

\*\*\* When the Engenuity version currently running on the storage system is more than 180 days older than the **Target Date**, **Days Old** displays the exact days old and the **Target Date** displays the target date.

## Converting directors

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

### Before you begin


This procedure requires Engenuity 5876.

### Procedure

1. Select a storage system.
2. Select **System > Hardware**.
3. Depending on the type of director you are converting, do the following in the **Hardware** panel:

Front end directors:

- a. Click **FE Directors** to open the Front End Directors list view.
- b.

Select a director, click  and then click **Convert FA to RF** to open the **Convert FA to RDF** dialog.

- c. Do 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to convert the director now.

RDF directors:

- a. Click **RDF Directors** to open the RDF Directors list view.
- b. Select a director and click **Convert RF to FA** to open the **Convert RDF to FA** dialog box.
- c. Do 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 920 and [Previewing jobs](#) on page 920.
  - Expand **Add to Job List**, and click **Run Now** to convert the director now.

4. (Optional) Open the Alerts list view (**Events > 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

### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. Click either the **FE Directors** or **RDF Directors** tab.
4. Select a port, click **Associate** to open the **Port Association** dialog.
5. Select a director and click **OK**.

## Setting director port attributes

### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. In the **Hardware** panel, click the type of director to open its list view.
- 4.

Select the director, click  and then click **Set Port Attributes** to open the **Set Port Attributes** dialog box.

5. (Optional) Select the port with the flag settings you want to copy.
6. Select or 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.

#### Show ACLX Volume

Enables the display of ACLX volume.

#### Common Serial Number

Enables multi-path configurations or hosts that need a unique serial number to determine which paths lead to the same volume.

#### Unique WWN

Ensures unique World Wide Names (WWNs) within the fiber environment (uses Symmetrix serial numbers and port numbers). This is enabled by default for all environment configuration changes and new environments. When disabled, you don't have to change WWNs.

#### Init Point to Point

Specifies a point-to-point (direct or switched) topology in the initialization sequence. When disabled (default), it is initialized as an arbitrated loop.

#### Volume Set Addressing

Enables the volume set addressing mode.

When using volume set addressing, you must specify a 4-digit address in the following range: (0)000-(0)007, (0)010-(0)017,... to a maximum of (0)FF0-(0)FF7.

The first digit must always be set to 0 (Symmetrix systems do not currently support the upper range of volume set addressing), the second digit is the VBus number, the third digit is the target, and the fourth digit is the LUN.

#### **Avoid Reset Broadcast**

Enables a SCSI bus reset to only occur to the port that received the reset (not broadcast to all channels).

#### **Negotiate Reset**

When enabled for AS/400 hosts, this flag forces a SCSI negotiation by the Symmetrix system after a SCSI reset, an error, or a bus volume reset.

#### **Enable Auto Negotiate**

Allows two fibre ports to handshake and settle on an optimal speed for data transfer.

#### **Environ Set**

Enables the environmental error reporting by the Symmetrix system to the host on the specific port.

#### **Disable Q Reset on UA**

When enabled, a Unit Attention (UA) that is propagated from another director does not flush the queue for this volume on this director. Used for hosts that do not expect the queue to be flushed on a 0629 sense (only on a hard reset).

#### **Soft Reset**

Supports SCSI soft reset on a Symmetrix port when enabled for a Bull/GCOS-7 host.

#### **SCSI 3**

Alters the inquiry data (when returned by any volume on the port) to report that the Symmetrix system supports the SCSI-3 protocol. When disabled, the SCSI-2 protocol is supported.

#### **SCSI Support1 (OS2007)**

Provides a stricter compliance with SCSI standards for managing volume identifiers, multi-port targets, unit attention reports, and the absence of a volume at LUN 0.

#### **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.

#### **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.

**OpenVMS**

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 920 and [Previewing jobs](#) on page 920.

## Associating directors and ports

**Before you begin**



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

When associating directors and ports, you can associate directors with ports, or ports with directors.

**Procedure**

1. Do one of the following:
  - To associate directors with ports:
    - a. Select the storage system.
    - b. Select **System > Hardware**.
    - c. In the **Hardware** panel, click **Available Ports** to open the **Available Ports** list view.
    - d. Select one or more ports.
    - e. Click **Associate** to open the **Associate Director** dialog box.
    - f. Select a director and click **Associate**.
  - To associate ports with directors:
    - a. Select the storage system.
    - b. Select **System > Hardware**.



- c. In the **Hardware** panel, click the director emulation type to open its list view.
- d. Select the director.
- e.
 
 Click  and then click **Associate** to open the **Port Association** dialog.
- f. Select the port and click **OK**.

## Disassociating directors and ports

### Before you begin

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

### Procedure

1. To disassociate ports with directors:
1. Select the storage system.
2. Select **System > Hardware**.
3. Click the tab for the appropriate director emulation type to open its list view.
4. Select the director.
5. Click **Disassociate**.
6. Click **OK**.

## Enabling and disabling director ports

### Procedure

1. Select the storage system.
2. Select **System > Hardware**.
3. Click either the **FE Directors** or **RDF Directors** tab.
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

### Before you begin

- To perform this procedure you must be an Administrator or Storage Admin.
- The storage system must be running Enginuity 5876 or higher.

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

To perform a health check on a storage system:

#### Procedure

1. Select a storage system.
2. Select the **System Health** tab.
3. In the **Actions** panel, click **Run Health Check**.


The Health Check view opens displaying results of previously run health checks, if any.

4. Click **Run Now** or **Add to Job List**.
5. Determine the health of the storage system by examining the **Status** and **Results** fields.

## Naming storage systems

This procedure explains how to assign custom *nice names* to storage systems:

#### Procedure

1. Select a storage system.
2. Select **System > Symmetrix Properties** (alternatively, click on **View Symmetrix Properties** in the **Actions** panel in the **System** view).
3. In the **Array Display Name** field, click on  .
4. 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.
5. Click the checkmark to apply the name or the X symbol to discard it.

## Replacing failed drives

#### 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.

---

**Note**

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.

---

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

## Step 1: Locate the failed drive

### Procedure

1. Select the storage system.
  2. Select the **System Health** tab.
  3. In the **Actions** panel, click **RUN DISK REPLACEMENT**.
  4. 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 .
- 

**Note**

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 ends, and you need to restart the process.

---

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

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


## Step 2: Replace the failed drive

### Before you begin


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.

### Procedure

1. To replace a failed disk drive:
  1. Access the disk drive. Do one the following, depending on the storage system model:
    - Storage Series system:
      - a. Open the front door of the storage 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.




- Storage SE system:
  - a. Open the front door of the storage system.
- storage 10K system:
  - a. Remove EMI front panel by pressing the side tabs and pulling straight forward.




CL4729


2. 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.



3. 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.
  - 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.



4. Reinstall the DAE cover or EMI front panel, or close the door, depending on the storage system model:
  - storage series system:
    - a. Press 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.
  - storage SE system:
    - a. Close the door.
  - storage 10K system:
    - a. 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

After you have completed the disk drive replacement, the lights will return to normal after several seconds. At this point, return to the Unisphere 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 to [Viewing the job list](#) on page 923 ):

**Spare Drive Replacement SUCCEEDED**  
If any other messages appear, follow the instructions.

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

## 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 configuration is a viable configuration for making changes.

Before creating new volumes, check for free physical disk space. New storage 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 storage 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 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 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 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 storage system.

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

**Table 136** Task status before and after server shutdown

Task status prior to server shutdown	Task status post server shutdown
Created	
Aborted	
Validated	
Running	
Successfully	
Run has error	
Prepare has error	
Submit has error	
Define has error	
Validate has error	
Done	
Failed	
Prepare in progress	Created
Abort in progress	
Submitted	
Prepared	
Defined	
Pending	

**Table 136** Task status before and after server shutdown (continued)

Task status prior to server shutdown	Task status post server shutdown
NA Acquiring lock Lock acquire failed Running During run During preview	
Commit in progress	Unknown <hr/> <b>Note</b> Configuration tasks that were in a Commit in progress state prior to server shutdown are 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 jobs

### Procedure

1. Select the storage system.
2. Select **Events > Job List** to open the Job List view.

## Scheduling jobs

This option can be used to schedule one-off jobs.

In addition, this option can be used to create a recurring daily SnapVX snapshot for a given time. There is also an option to cancel a recurring snapshot. In the event of a failed recurring snapshot, an Alert will be raised to notify the user. The schedule continues to run in the event of a failed snapshot, issuing alerts to the user. The alerts list view retains a record of the failed snapshots (unless the alert is deleted). A warning level alert is issued.

Note the following for SnapVX snapshot jobs:



- If the job is already scheduled to be recurring, then the job cannot be grouped.
- Unscheduled jobs can be grouped together and then a single recurring schedule can be applied to that group.
- You cannot edit a recurring schedule once it has been created. If you want to change a recurring schedule, you must delete the existing recurring schedule and set up a new recurring schedule.
- If you want to create more than one daily snapshot, for example, every Monday at 9AM create a snapshot and every Monday at 5PM create a snapshot, you have to



schedule multiple jobs for the desired times, that is, re-run the Protection Wizard with SnapVX and add to job list and schedule the jobs for the additional times.

- When scheduling a weekly snapshot, it is recommended that you choose one day of the week and time to take the snapshot, for example, selecting Wednesday at 8AM in order for the snapshot to occur every Wednesday at 8AM.

#### Procedure

1. Select the storage system.
2. Select **Events** > **Job List** to open the Job List view.
3. Select a job and click **Schedule** to open the **Schedule** dialog box.
4. Select an **Execution Time** and an **Execution Date**.
5. SnapVX Hourly snapshots only: Do one of the following: Select the **Occurrence** option—Daily, Hourly, or Run Once.
  - Daily: Select the **Execution Time** and select the days you want the job to reoccur on at the selected time.
  - Hourly: Select an **Execution Time** and an **Execution Date**.
  - Run Once: Select an **Execution Time** and an **Execution Date**.
6. Click **OK**.
7. Optional: Deleting a scheduled job
  -  Select the job, click  and then click **Delete**.
  - Click **OK** to confirm.

## Running jobs

#### Procedure

1. Select the storage system.
2. Select **Events** > **Job List** to open the Job List view.
3. Select a job and click **Run**.
4. Click **OK**.

## Rescheduling jobs

#### Before you begin

You cannot reschedule a single task in a job, only the entire job.

If the volumes involved in a configuration change were originally reserved, rescheduling the tasks without reservation does not release the reservations. You need to manually release the reservations.

To reschedule a job:

#### Procedure

1. Select the storage system.
2. Select **Events** > **Job List**.
3. Select a job and click **Schedule** to open the **Schedule** dialog box.

4. Select a new **Execution Date** and **Execution Time**.
5. Click **OK**.

## Modifying jobs


### Procedure

1. Select the storage system.
2. Select **Events** > **Job List** to open the Job List view.
3. Select a job and click **Modify**.
4. Modify the name, schedule or the occurrence and click **OK**.

Note: Job names must be unique from other jobs on the storage system.

## Reordering tasks within a job

### Procedure

1. Select the storage system.
2. Select **Events** > **Job List** to open the Job List view.
3. Select a job and click , then click on the link in the **Status** field.
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.

### Procedure


1. Select the storage system.
2. Select **Events** > **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.
5. Use the calendar drop-down menu to schedule the date and time of the job.
6. Click **OK**.

## Un-grouping jobs

When un-grouping a job, Unisphere creates an individual job for each of the tasks in the original job.

### Procedure

1. Select the storage system.
2. Select **Events** > **Job List** to open the Job List view.
- 3.

Select a job and click .

4. In the task list, select **Ungroup** for each of the tasks you want to un-group from the job.

## Stopping jobs

### Procedure

1. Select the storage system.
2. Select **Events > Job List** to open the Job List view.
3. Select an active job and click **Stop**.

If Unisphere successfully stops the job, the state of the job changes to **Stopped**.

## Deleting jobs

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

### Procedure

1. Select the storage system.
2. Select **Events > Job List** to open the Job List view.
- 3.

Select one or more jobs, click  and then click **Delete**.

4. Click **OK**.

Unisphere removes the jobs from the job list view.

## Viewing the job list

### Procedure

1. Select the storage system.
2. Select **Events > 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 assigned ID, and an arrow icon indicating if the job contains one or more tasks.

#### Status

Status of the job. Possible values are:

##### SCHEDULED

Job is scheduled for execution.

##### UNSCHEDULED

Job is not 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.

#### ABORTED

Job was running and a user aborted it.

#### 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 jobs.

#### Recurring


Indicates whether the job is recurring.

The following controls are available:

- **Modify**—[Modifying jobs](#) on page 922
- **Run**—[Running jobs](#) on page 921
- **Stop**—[Stopping jobs](#) on page 923
- **Schedule**—[Scheduling jobs](#) on page 920
- **Group**—[Grouping jobs](#) on page 922
- **Ungroup**—[Un-grouping jobs](#) on page 922
- **Delete** — [Deleting jobs](#) on page 923
-  — [Viewing job details](#) on page 924

## Viewing job details

### Procedure

1. Select the storage system.
2. Select **Events** > **Job List** to open the Job List view.
3. Select a job and click .

The following properties are displayed:

#### Status

Status of the job. Possible values are:

#### SCHEDULED

Job is scheduled for execution.

**UNSCHEDULED**

Job is not scheduled to run.

**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.

**SUCCEEDED**

Job succeeded.

**ABORTED**

Job was running/appending and a user aborted it.

**Owner**

Host from which the job was created and the ID of user who created it.

**Scheduled****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 jobs.

## Understanding licenses

Unisphere supports 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 customer service, copy them to a Solutions Enabler or a Unisphere host, and load them onto 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, obtain a new license file from online support and load it onto the storage system.

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

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

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

Unisphere uses array-based eLicensing.

---

#### Note

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

---

## Installing licenses



### Before you begin

- To perform this operation, you must be an Administrator or StorageAdmin.
- Before you can use Unisphere to manage a storage system, you must first install a license on the storagesystem.
- To obtain array-based licenses from the support website you will need the License Authorization Code (LAC) identification number from the LAC letter emailed to you.

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

To install licenses:

### Procedure

1. Select the storage system.
2.  Click  in the title bar to open the **Settings** dialog, then click **System and Licences**.
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 support 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.
  - d. Click **Choose File** and select the license file.  
The selected file's content loads for you to preview before loading it on the storage system.
  - e. Click **OK** to load the license file on the storage 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

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **System and Licences** > **Solutions Enabler** to open the **Solutions Enabler** license view.
3. Select the license and click .

## Viewing Symmetrix entitlements

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **System and Licences** > **Symmetrix Entitlements** to open the Symmetrix Entitlements list view.

Use the **Symmetrix Entitlement** list to view all licensed features (including those licensed with array-based entitlements and host-based keys, which are still required) on storage systems running Enginuity 5876 or higher. In addition, you can obtain and install license files on storage systems running Enginuity 5876 or higher.

The License Usage information is displayed for the first of the listed storage arrays. You can select a storage system from the dropdown list to view its license usage information.

The following properties display:

#### Feature Name

Name of the feature.

#### License

Whether the license is host-based (SE) or array-based (EMCLM).



#### Status

Whether the license is Enabled or Disabled.

#### Expiration Date

Date an evaluation license expires. For permanent licenses, this field appears blank.

The following controls are available:


-  —To view details of a license file installed on the storage system, hover over the entry and click  (see [Viewing license file](#) on page 929).

- **Get File**—Opens the support website, from which you can obtain a license file (by downloading or through email). Download or copy a license file to the SMAS server before applying it to a storage system.
- **Load File**—[Installing licenses](#) on page 926

## Viewing host-based licenses

### Before you begin

#### Procedure

1. Click  in the title bar to open the **Settings** dialog, then click **System and Licences** > **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**—[Installing licenses](#) on page 926
-  — [Removing host-based licenses](#) on page 927

## Viewing license usage

### Procedure

1. Click  in the title bar to open the **Settings** dialog, then click **System and Licences** > **License Usage** to open the license usage list view.

The License Usage view shows the results of a query to the storage system's feature registration database (Enginuity 5876 or higher).

The License Usage information is displayed for the first of the listed storage arrays. You can select a storage system from the dropdown list to view its license usage information.

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 customer service.

Product titles activated manually (MAN) or because they were in use (USE) are not considered properly entitled, in which case contact customer service for proper entitlement.

**Capacity Type**

Qualifies the licensed capacity. Valid 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 storage system.

**R-TB External**

Indicates that the capacity licensed applies to the raw capacity of the virtualized LUNs in external storage.

**Licensed Capacity (TB)**


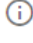
Maximum quantity of data for which the functionality of the software is licensed to use, in Terabytes.

**Used Capacity (TB)**

Amount of licensed capacity currently in use.

## Viewing license file




**Procedure**

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **System and Licences > Symmetrix Entitlements** to open the Symmetrix Entitlements list view.  
  
Use the **Symmetrix Entitlement** list to view all licensed features (including those licensed with array-based entitlements and host-based keys, which are still required) on storage systems running Enginuity 5876 or higher. In addition, you can obtain and install license files on storage systems running Enginuity 5876 or higher.  
  
The License Usage information is displayed for the first of the listed storage arrays. You can select a storage system from the dropdown list to view its license usage information.
3. Hover over the row whose license file you wish to see and click .  
The following properties are displayed:

- **Feature Name**—Name of the feature.
  - **License**—Whether the license is host-based (SE) or array-based (EMCLM).
  - **Status**—Whether the license is Enabled or Disabled.
  - **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 storage systems or to all the storage systems in the enterprise. This column is blank for Evaluation type licenses.
  - **Install Date**—Date the license file was installed on the storage system.
  - **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 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 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 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 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.
4. Click **VIEW FILE** to view the formatted file or the raw output.

## Viewing license file details

### Procedure

1. Select the storage system.
2.  Click  in the title bar to open the **Settings** dialog, then click **System and Licences > Symmetrix Entitlements** to open the Symmetrix Entitlements list view.
3. Hover over the row whose license file you wish to see and click .
4. Click **VIEW FILE** to view the formatted file or the raw output.  
A dialog with two tabs (**Formatted File** and **Raw Output**) is displayed.
5. Click **Formatted File** to view the license file as a formatted file
6. Click **Raw Output** to view the raw output of the license file.

## 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.

---

### Note

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

---

To set up access controls:

### Procedure

1. Create an access control group. See [Creating access groups](#) on page 932.
2. Add to the group, one or more hosts (access name) and unique IDs.  
Each host has a unique ID; a group can contain one or more hosts. See [Adding access ID to access groups](#) on page 932.
3. Create one or more pools of volumes.  
Specific volumes can belong to only one pool. See [Creating access groups](#) on page 932.
4. Create one or more access control entries.  
An access control entry associates a pool with a group, and grants the access control types. See [Creating access control entries](#) on page 937.

## Opening access controls



### Before you begin

Access control dialogs and views are restricted and require you to enter a PIN provided by customer service.

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

To open access controls:

### Procedure

1. Select a storage system.
2.  Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control** and select **Access Control Entries**, **Access Groups** or **Access Pools**.
3. To enable active management, click on the PIN link in the read-only warning or else on the lock icon next to the section title.
4. In the **Enter PIN** dialog, enter your Access Control PIN and click **OK**.  
Active management is now enabled and the read-only warning no longer appears.


## Creating access groups

Typically various sets of users tend to use the same applications that utilize common 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:

### Procedure

1. Select a storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Groups**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Click **Create** to open the **Create Access Group** dialog box.
5. Type a **Group Name**.  
Access group names must be unique from other access pools on the system and cannot exceed 31 characters (Engenuity 5876 or higher). Only alphanumeric characters, underscores ( \_ ), and dashes ( - ) are allowed. Access group names are case-sensitive.
6. Type the **Host ID**.  
This value is the name of the access ID (eight characters).
7. 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).
8. Click **Add** to add an access ID to the access group (see [Adding access ID to access groups](#) on page 932).
9. Click **OK**.

## Adding access ID to access groups



### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Groups**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select the group and click .

5. In the **Properties** panel, click the number next to **Access IDs**.
6. Click **Add** to open the **Add Access ID** dialog box.
7. Do one of the following:
  - 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. Example host ID: 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**.
8. Click **OK**.


## Removing access IDs from access groups

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Groups**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select the group and click .
5. In the **Properties** panel, click the number next to **Access IDs** to open the Access ID list view.
6. Select one or more IDs and click **Remove**.
7. Click **OK**.

## Deleting access groups

### Procedure


1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Groups**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select one or more groups and click **Delete**.
5. Click **OK**.

## Viewing access groups

### Before you begin

A read only user can view access groups. You must enter a PIN in order to make changes.

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Groups**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. 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:

-  — [Viewing access group details](#) on page 934
- **Create** — [Creating access groups](#) on page 932
-  — [Deleting access groups](#) on page 933

## Viewing access group details

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Groups**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select the group and click .
5. Use the access group details view to display and manage an access group. This view contains the following panel:

#### 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 **Properties** panel provides links to views for objects contained in and associated with the access group. For example, clicking **Access IDs** opens a view listing the access IDs in the group (see [Viewing access IDs](#) on page 939).

## Creating access pools

**Procedure**

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Pools**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Click **Create** to open the **Create Access Pool** dialog box.
5. Type a **Pool Name**.

Access pool names must be unique from other access pools on the system and cannot exceed 31 characters (Engenuity 5876 or higher). Only alphanumeric characters, underscores ( \_ ), and dashes (-) are allowed. Access pool names are case-sensitive.

6.



Select one or more available volumes and click to select it.

**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.


7. Click **OK**.

## Modifying access pools

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


**Procedure**

1. Select a storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Pools**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931

4. Select the access pool and click  to open its details view.
5. In the **Properties** panel, click the number next to **Volumes** to the access pool volumes list view.
  - 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**.
    - c. Click **OK**.
  - To remove volumes from the pool:
    - a. Select one or more volumes and click **Remove**.
    - b. Click **OK**.

## Deleting access pools

### Procedure


1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Pools**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select one or more pools and click **Delete**.
5. Click **OK**.

## Viewing access pools

### Before you begin

A read only user can view access pools. You must enter a PIN in order to make changes.

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Pools**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Use the Access Pools list view to display and manage access pools on the storage system. Select a storage system ID from the drop-down list.

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 control entries.





The following controls are available:

-  — [Viewing access pool details](#) on page 937
- **Create**—[Creating access pools](#) on page 935
-  — [Deleting access pools](#) on page 936

To view access pool volumes, see [Viewing access pool volumes](#) on page 940.

## Viewing access pool details

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Pools**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select the access pool and click  to open its details view.
5. Use the Access Pool details view to display and manage an access pool.

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**—[Creating access pools](#) on page 935
-  — [Deleting access pools](#) on page 936


The **Properties** panel provides links to views for objects contained in and associated with the access pool. For example, clicking the number next to **Volumes** opens a view listing the volumes contained in the pool.

## Creating access control entries

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Control Entries**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931


4. Click **Create**.
5. Select/Create the access group to include in the ACE.
6. Select/Create the access pool to include in the ACE.
- 7.

Select one or more available access types and click  to select it.

8. Click **OK**.

## Deleting access control entries

### Procedure


1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Control Entries**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select one or more entries and click **Delete**.
5. Click **OK**.

## Viewing access control entries

### Before you begin

A read only user can view access control entries. You must enter a PIN in order to make changes.

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Control Entries**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Use the Access Control Entries list view to display and manage access control entries on the storage system. Select a storage system ID from the drop-down list.

The following properties display:

#### Access ID Group

Associated access group.



#### Access-controlled Volume Pool

Associated access pool.

#### Access Types



Permissions assigned to the group/pool. For information about valid values, refer to [Access types](#) on page 941, [Viewing access types](#) on page 940, and [Modifying access types](#) on page 942.

The following controls are available:

-  — [Viewing access control entry details](#) on page 939
- **Create** — [Creating access control entries](#) on page 937
-  — [Deleting access control entries](#) on page 938

## Viewing access control entry details

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Control Entries**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select the entry and click .
5. Use the details view to list and manage access control entries.

### 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 [Access types](#) on page 941.

The **Properties** panel provides links to views for objects contained in and associated with the access control entry. For example, clicking **Access IDs** opens a view listing the access IDs associated with the entry.

## Viewing access IDs

### Procedure

1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Groups**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select the group and click .
5. In the **Properties** panel, click the number next to **Access IDs**.
6. Use the Access ID list view to display and manage access IDs.  
The following property displays:



**Access ID**

Access IDs assigned to the group.

The following controls are available:

- **Add**—[Adding access ID to access groups](#) on page 932
- **Remove**—[Removing access IDs from access groups](#) on page 933

**Viewing access pool volumes****Procedure**


1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Pools**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select the access pool and click  to open its details view.
5. In the **Properties** panel, click the number next to **Volumes**.
6. Use the Access Pool Volumes list view to display and manage the volumes in an access pool.

The following property displays:



**Volume ID**

Volume identifier.

The following controls are available:

- **Add**—[Modifying access pools](#) on page 935
-  —[Modifying access pools](#) on page 935

**Viewing access types****Procedure**


1. Select the storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Control Entries**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4. Select the entry and click  .
5. In the **Properties** panel, click the number next to **Access Types**.
6. Use the Access Types list view to display and manage the access types associated with an access control entry.

The following property displays:

### Access Type

Permissions assigned to the group/pool. For more information about valid values, refer to [Access types](#) on page 941 .

The following controls are available:

- **Add**—[Adding access ID to access groups](#) on page 932
-  —[Removing access IDs from access groups](#) on page 933

## Access types

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

**Table 137** Access types

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 storage system.
CFGSYM	Allows access to set storage system attributes, set port flags, and swap RA groups. Must be directed to ALL volumes.
CHECKSUM	Allows volume Double Checksum operations.
CREATEDV	Allows the creation and deletion of volumes.
DIRCTRL	Allows you to take directors and their ports offline and online. Must be directed to ALL volumes.





**Table 137** Access types (continued)

Access type	Description
ECC	Allows the ECC 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.
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.

## Modifying access types

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

### Procedure

1. Select the storage system.
2.  Click  in the title bar to open the **Settings** dialog, then click **Symmetrix Access Control > Access Control Entries**.
3. To enable active management, enter your Access Control PIN: [Opening access controls](#) on page 931
4.  Select the entry and click .
5. In the **Properties** panel, click the number next to **Access Types** to open the Access Types list view.
  - To add access types:
    - a. Click **Add** to open the **Assign Access Types** dialog box.
    - b. Select one or more access types and click **Add**. For more on the available access types, see [Access types](#) on page 941

- c. Click **OK**.
- To remove access types:
  - a. Select one or more access types and click **Remove**.
  - b. Click **OK**.

## 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 specified donation time.

---

### Note

Engenuity 5876 is required for actively managing dynamic cache partitions. DCPs can be viewed on storage systems running HYPERMAX OS 5977 Q316SR or higher but they can't be actively managed.


---

## Enabling/Disabling dynamic cache partitioning

### Before you begin

This feature is not supported on HYPERMAX OS 5977 or higher.

### Procedure

1. Select a storage system.
2. Click  in the title bar to open the **Settings** dialog, then click **Management > Symmetrix Attributes**.
3. In the **DCP** panel, set the **Cache Partition Status** to **Enable** or **Disable**.
4. Click **Apply**.

## Creating dynamic cache partitions

### Before you begin

- This feature is not supported on HYPERMAX OS 5977 or higher.
- 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%.

### Procedure

1. Select the storage system.

2. Select the **System Health** tab.
3. In the **Actions** panel, click **View Other Hardware** and then **Cache Partitions**.
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 31 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 value 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**.

## Modifying dynamic cache partitions

### Before you begin

This feature is not supported on HYPERMAX OS 5977 or higher.

### Procedure

1. Select the storage system.
2. Select the **System Health** tab.
3. In the **Actions** panel, click **View Other Hardware** and then **Cache Partitions**.
4. Select a dynamic cache partition and click **Modify**.
5. Modify the **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 31 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. Modify the minimum target percentage (**Min Target %**) for the partition. This value must be less than the Target %.
7. Modify the target cache percentage (**Target %**) for the partition. This value must be less than the Max Target %.
8. Modify the maximum cache percentage (**Max Target %**) for the partition.
9. Modify 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. Modify 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.

### Procedure

1. Select the storage system.
2. Select **Storage > Volumes**.
3. Click on the appropriate volume panel.
- 4.

Select the volume, click  , and select **Assign Dynamic Cache Partition**.

5. Select a **Dynamic Cache Partition** and click **OK**.
6. Click **OK**.

## Assigning dynamic cache partitions

### Before you begin

This feature is supported on storage systems running Enginuity 5876 and on storage systems running HyperMax OS 5977 Q316SR and higher.

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.

### Procedure

1. Select the storage system.
2. Select **STORAGE > Volumes**.
3. Navigate to the volume.
- 4.

Select a volume, click  , and click **Assign Dynamic Cache Partition**.

5. Select a dynamic cache partition name and click **OK**.

## Deleting dynamic cache partitions

### Before you begin

This feature is not supported on HYPERMAX OS 5977 or higher.

### Procedure

1. Select the storage system.
2. Select the **System Health** tab.
3. In the **Actions** panel, click **View Other Hardware**.

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**


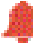



Number of cache slots used by the partition.

**% 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**—[Creating dynamic cache partitions](#) on page 943
- **Delete**—[Deleting dynamic cache partitions](#) on page 945
-  —[Viewing dynamic cache partitions](#) on page 947

4. Select a partition 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.

#### Procedure

1. Enable dynamic cache partitioning in Analyze mode under **Settings > Symmetrix Attributes > DCP**.
2. Create your dynamic cache partitions. See [Creating dynamic cache partitions](#) on page 943.
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

#### Procedure

1. Select the storage system.
2. Select the **System Health** tab.
3. In the **Actions** panel, click **View Other Hardware**.

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**—[Creating dynamic cache partitions](#) on page 943
- **Delete**—[Deleting dynamic cache partitions](#) on page 945
-  —[Viewing dynamic cache partitions](#) on page 947

## Viewing dynamic cache partition details

### Procedure

1. Select the storage system.
2. Select the **System Health** tab.
3. In the **Actions** panel, click **View Other Hardware**.

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**—[Creating dynamic cache partitions](#) on page 943
- **Delete**—[Deleting dynamic cache partitions](#) on page 945
-  —[Viewing dynamic cache partitions](#) on page 947

4.

Select a partition and click .

The Cache Partition details view allows you to view and manage a cache partition. This view contains **Details** and **Performance** panels.

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.


**Volumes**

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

This panel displays with inactive links if the selected storage system is not registered for data collection.

## Viewing volumes assigned to dynamic cache partitions

**Procedure**

1. Select the storage system.
2. In the **Dashboard**, click the **System Health** tab.
3. In the **Actions** panel, click **View Other Hardware**.
4. In the **Cache Partitions** tab, select a partition and click  to open it **Details** view.
5. Click on the number in the **Volumes** field to open the **Volumes** list view.

Use the **Volumes** list view to display and manage volumes assigned to the partition.

The following properties display:

- **Name**—Assigned volume name.
- **Type** — Type of volume.
- **Meta Config**—Volume configuration.
- **Striped Size**—Meta striped size.
- **Status**—Volume status.
- **Reserved**—Indicates whether the volume is reserved.
- **Capacity (GB)**—Volume capacity in Gigabytes.
- **Emulation**—Emulation type for the volume.
- **Paths**—Number of masking records for the volume.

## System management - iSCSI

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

iSCSI is a protocol that uses the TCP to transport SCSI commands, enabling the use of the existing TCP/IP networking infrastructure as a SAN. As with SCSI over Fibre Channel (FC), iSCSI presents SCSI targets and devices to iSCSI initiators (requesters). Unlike NAS, which presents devices at the file level, iSCSI makes block devices available 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.

To view the iSCSI dashboard

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** 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.
- **Actions**—A panel with the following controls:
  - **Create iSCSI Target**—[Creating an iSCSI target](#) on page 952
  - **Create IP Interface**—[Creating an IP interface](#) on page 954
  - **Add IP Route**—[Adding an IP route](#) on page 955
- **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.

**Description**

The description of the alert.

**Created**

Date/time the alert was created.

**Acknowledged**

The date and time that the alert was acknowledged.

## Creating an iSCSI target

**Procedure**

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** to open the iSCSI dashboard.
3. In the **Actions** panel, click **Create iSCSI Target**.
4. Select a director from the drop-down menu.
5. (Optional) (Optional) Select the **Use custom name** check box and then 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.

6. Type a value for the **Network ID**.  
Valid values range from 1 through 16383.
7. Type a value for the **TCP Port**.  
Default value is 3260. Valid values range from 0 through 65535.
8. (Optional) Click **Advanced Options** 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 ISID and iSCSI name to calculate the World Wide Name (WWN) in the Persistent Group Registration (PGR) record.

9. Click **OK**.

## Modifying an iSCSI target

**Procedure**

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** to open the iSCSI dashboard.
3. Click **iSCSI Targets** to open the **Attached iSCSI Targets** list view.
4. Select a director and click **Modify** to open the **Modify iSCSI Target** dialog.
5. (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.

6. Type a value for the **Network ID**.

Valid values range from 1 through 16383.

7. Type a value for the **TCP Port**.

Default value is 3260. Valid values range from 0 through 65535.

8. Click **OK**.

## Creating an IP interface

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** to open the iSCSI dashboard.
3. In the **Actions** panel, click **Create IP Interface**.
4. Select a director and port combination from the drop-down menu.
5. 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.
6. 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.
7. Type the value for the **Network ID**.  
Valid values range from 1 through 16383.
8. Type a value for the **VLAN ID**.  
Valid values range from 0 through 4094. This setting fails if same VLAN id is used for more than one IP interface on a specified SE physical port.
9. (Optional) Type a value for the **Maximum Transmission Unit**.  
Default value is 1500. Valid values range from 1500 through 9000.
10. Click **OK**.

## Editing an IP interface

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** to open the iSCSI dashboard.
3. Click **IP Interfaces**.
4. Select an IP interface and click **Modify**.
5. Select a director and port combination from the drop-down menu.
6. Modify the 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. Modify the 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. Modify the value for the **Network ID**.  
Valid values range from 1 through 16383.

9. (Optional) Modify the value for the **Maximum Transmission Unit**.  
Default value is 1500. Valid values range from 1500 through 9000.
10. Click **OK**.

## Adding an IP route

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** to open the iSCSI dashboard.
3. In the **Actions** panel, click **Add IP Route**.
4. Select a director and port combination from the drop-down menu.
5. Type a value for the **Destination IP**.  
You specify the IPv4 address in dotted decimal notation. You specify the IPv6 address in colon-hexadecimal format.
6. 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.
7. 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.
8. (Optional) Type the value for the **Network ID**.  
Default value is 1. Valid values range from 1 through 16383.
9. Click **OK**.

## Deleting an iSCSI target

Deleting iSCSI targets with IP interfaces attached detaches the IP interfaces before deleting the iSCSI target.

The deletion fails if the target is in a port group.

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** to open the iSCSI dashboard.
3. Click **iSCSI Targets** (Attached or Unattached).
4. Select an iSCSI target and click **Delete**.
5. Click **OK**.

## Deleting an IP interface

---

### Note

An IP interface attached to an iSCSI target cannot be deleted.

---

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.

2. Select **System** > **iSCSI** 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 **Yes** to confirm the operation.

## Removing an IP route

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** 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 **Yes** to confirm the operation.

## Attaching an IP interface to an iSCSI target

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** to open the iSCSI dashboard.
3. Click **IP Interface**.
4. Select an IP interface and click **Attach**.
5. Select whether to attach to an existing or new iSCSI target, and do one of the following depending on your selection:
  - Existing:
    - a. 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**.  
The default value is 3260. Valid values range from 0 through 65535.
    - c. Optional: Click **Advanced Options** 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 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 storage 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**.

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

**Procedure**

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** 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 one of the following depending on your selection:

- Existing:
  - a. Select an IP interface from the list.  
You can select 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**.  
The 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**.  
If you toggle between the **New** and **Existing** selections, the system applies only the changes from within your last selection.
- 7. Click **OK** to confirm your changes.

## Detaching an IP interface from an iSCSI target

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** 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**.

## Disabling an iSCSI target

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** 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

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** 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

### Before you begin

The iSCSI target must be in an offline state before the flags can be modified.

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** to open the iSCSI dashboard.
3. Click **iSCSI Targets** (Attached or Unattached).
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 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 storage 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 ISID 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

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** 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.

To see more information on a director, select it and click 




The following controls are available, depending on the operating environment:

- **Create iSCSI Target**—[Creating an iSCSI target](#) on page 952
- **Add IP Route**—[Adding an IP route](#) on page 955

## Viewing the iSCSI director details

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** 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  to open the iSCSI Directors details view.

The following panels and properties display, depending on the operating environment:

Properties panel:

#### Director

The name of the iSCSI director.

#### Ports

The number of the 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 iSCSI targets attached to the IP interfaces associated with the director.

#### Unattached iSCSI targets

The total number of 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 port.

## Viewing IP interfaces list

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System > iSCSI** 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:

-  — [Viewing IP interfaces details](#) on page 962
- **Create**—[Creating an IP interface](#) on page 954
- **Attach**—[Attaching an IP interface to an iSCSI target](#) on page 956
- **Detach**—[Detaching an IP interface from an iSCSI target](#) on page 958
- **Delete**—[Deleting an IP interface](#) on page 955

## Viewing IP interfaces details

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** 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  to open the IP Interfaces details view.

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**

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(OS2007)**

Provides a stricter compliance with SCSI standards for managing volume identifiers, multi-port targets, unit attention reports, and the absence of a volume at LUN 0.

### **SPC2 Protocol Version**

#### **Open VMS**

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.

#### **ISID Protected**

Enables an Open VMS fiber connection.

#### **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 iSCSI targets list**

### **Procedure**

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** 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:

-  — [Viewing iSCSI target details](#) on page 965
- **Create** — [Creating an iSCSI target](#) on page 952
- **Enable** — [Enabling an iSCSI target](#) on page 959
- **Disable** — [Disabling an iSCSI target](#) on page 958
- **Set Port Flags** — [Setting port flags](#) on page 959
- **Attach** — [Attaching an IP interface to an iSCSI target](#) on page 956
- **Modify** — [Modifying an iSCSI target](#) on page 953
- **Detach** — [Detaching an IP interface from an iSCSI target](#) on page 958
- **Delete**  — [Deleting an iSCSI target](#) on page 955

## Viewing iSCSI target details

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** 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  to open the **iSCSI Targets** details view.

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 Port

The TCP port to be used for all IP addresses attached to it (editable).

#### Status

The status of the iSCSI target.

#### IP Interfaces

The IP Interfaces attached to the iSCSI target.

#### iSCSI Ports

The physical iSCSI ports associated with the iSCSI target.

#### Volumes

The total number of volumes mapped to the iSCSI target.

#### 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 IP routes list****Procedure**

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** 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 with which the IP route is associated.

The following controls are available, depending on the operating environment:

-  — [Viewing the IP routes details](#) on page 966
- **Create**—[Adding an IP route](#) on page 955
- **Remove**—[Removing an IP route](#) on page 956

**Viewing the IP routes details****Procedure**

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** 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  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 with which the IP route is associated.

The following control is available, depending on the operating environment:

- **Remove**—[Removing an IP route](#) on page 956

## Viewing iSCSI ports list

**Procedure**

1. Select a storage system running or higher.
2. Select **System** > **iSCSI** 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 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**—[Creating an IP interface](#) on page 954
-  —[Viewing iSCSI ports details](#) on page 968

## Viewing iSCSI ports details

### Procedure

1. Select a storage system running HYPERMAX OS 5977 or higher.
2. Select **System** > **iSCSI** 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  to open the **iSCSI Ports** details view.

### Properties panel

The following properties display, depending on the operating environment:

#### Director

The name of the director and the iSCSI port combined.

#### Status

The current status of the iSCSI port (offline or online).

#### IP Interfaces

The total number of IP interfaces attached to this iSCSI port.

#### Attached iSCSI Targets

The list view of iSCSI targets on this director that are attached to IP interfaces.

#### Volumes

The total number of volumes mapped to all iSCSI targets on this port.

#### Speed GB/Sec

The speed of the iSCSI port.

#### 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.