

Dell EMC Solutions Enabler Array Controls and Management

CLI User Guide

9.0

Notes, cautions, and warnings

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

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

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

Figures.....	16
Tables.....	17
Please Read: Solutions Enabler and Engenuity 5876 document support.....	19
PREFACE.....	20
Part I: Setting up and Configuring Storage Arrays.....	23
Chapter 1: Introduction	24
Solutions Enabler overview.....	24
SYMCLI help options.....	24
Environment variables.....	25
Command input and output tips.....	27
Version compatibility information and restrictions.....	27
Capacity information.....	28
Chapter 2: Discovery.....	30
Discovery overview.....	30
Discovery syntax options.....	30
Verify host configuration database is synchronized	31
Scan for new devices	31
Connectivity authorization.....	32
Configuring virtual disk mapping.....	32
Display virtual disks as array devices.....	33
Host configuration database file.....	33
Database file location	33
Database file lock.....	34
Changing the current database file name.....	34
Database location in client/server mode	34
Database access modes	34
Command modes: Online and offline	34
Database configuration information	35
Configuration data synchronization.....	35
Update configuration information (sync operation)	35
Inhibit Database synchronization	36
Chapter 3: User and Group Authorization.....	37
User and group authorization overview.....	37
User roles	37
Assign multiple user roles	38
User authorization rules	38
User ID formats	38
Authorization settings management.....	40

Current username identification.....	40
Manage Host IDs	40
User authorization	42
User authorization enforcement	43
Set Secure Reads.....	44
Modify user authorization role mappings	44
User authorization command file usage	45
Preview and commit user authorization actions.....	45
Backup the authorization database	46
Restore the authorization database.....	46
Monitor authorization	46
List user roles, names, and components.....	47

Chapter 4: Host-based Access Control..... 49

Host-based access control overview.....	49
Access Control database.....	49
Host access IDs.....	50
Create or modify access control data	53
Applying access control operations.....	53
Minimum access control configuration	54
Initial ACL setup	54
Create and manage access groups.....	55
Verify administrative authority	55
Create an access group.....	55
Add host access IDs to a group.....	55
User access IDs (PINs) for AdminGrp	56
Edit and manage an access group.....	56
Create and manage limited access to access pools.....	58
Verify administrative authority	58
Create an access pool.....	58
Add devices to access pool	59
Edit and manage access pools	59
Create and manage access control entries	60
Verify administrative authority	60
Grant permissions to access group.....	60
Available access control permissions	61
Grant BASE permissions to a group.....	63
Grant SRDF permissions to a group.....	64
Grant BCV and SDR permissions to a group.....	64
Grant BASE permissions to a non-registered host.....	64
Remove permissions from access group.....	64
Obtain access control information	65
List access control information	65
Show access control information	65
Obtain host access ID	65
Verify a locked access control session.....	66
Release locked access control sessions	66
Access control strategies.....	66
Default configuration: all permissions to users and hosts.....	66
Manage default access	67

Setting up access control for TimeFinder devices.....	69
Setting access control for SRDF devices.....	69
Unisphere for VMAX setup strategy	69
Setup access control for Symaccess.....	69
Backup and restore the access control database	69
Create an access control database backup file	69
Restore access control database with backup file	70
Chapter 5: Thin Device Management	71
Thin device management and reporting overview.....	71
Allocate space on thin devices.....	71
Reclaim allocations on thin devices.....	72
Free allocations or free allocations with written tracks	73
Background operations on thin devices.....	73
Stop background operations on thin devices.....	73
Rename thin pools.....	74
Verify pool and device states.....	74
Monitor thin pools.....	74
View all thin device pools.....	75
View thin device pool details.....	76
View thin devices.....	77
View thin device details.....	78
View thin device allocations bound to a different pool.....	79
Chapter 6: Grouping Devices.....	81
Overview of groups.....	81
Device groups.....	81
Group name services	81
Names of device groups and devices	81
Device group types	82
Device lists	82
Device restrictions	82
Create a device group	84
Add devices to a device group	84
Set controls on Celerra devices.....	87
List devices in a device group.....	87
List all device groups.....	88
Show device group details	88
Export and import device lists	89
Rename device groups	90
Rename logical device names.....	91
Move a device between device groups.....	91
Move all devices between device groups.....	92
Copy devices between device groups	92
Copy devices from device group to storage group	93
Remove a device from a device group.....	93
Remove all devices from a device group.....	94
Delete a device group.....	94
Perform control operations on device group devices	95

Storage groups.....	95
Create storage groups.....	95
Create devices and add to storage groups.....	97
Add existing devices to storage groups	98
Add child storage groups to existing storage groups	99
Storage group merge operation.....	99
Storage group split operation.....	99
Modify storage group properties	100
Standalone storage groups and cascaded groups	104
Remove devices from a storage group	106
Remove child storage groups	107
Host I/O Limits for storage groups	107
Copy devices from a storage group to a device group	110
List storage groups	110
View storage group details	113
Show cascaded storage group details.....	113
Export and import device lists	115
Rename a storage group.....	117
Move devices between storage groups.....	117
Copy a device between storage groups	119
Copy all devices between storage groups.....	119
Copy devices from a storage group to a device group	120
Delete a storage group	121
Perform control operations on storage group devices	121
Composite groups.....	121
Composite group device members	121
Logical device name support.....	122
Device group members within composite groups.....	122
Logical names of devices within a composite group	122
Composite group types.....	123
SRDF consistency groups.....	123
Create a composite group	123
Export a composite group to file.....	125
Delete a composite group	126
Import a composite group.....	126
Add standard devices to a composite group.....	126
Move and copy devices in composite groups.....	127
Remove a standard device from a composite group.....	129
Rename a composite group.....	130
Set SRDF group attributes.....	130
Composite group creation and output	131
Display composite group information	132
Show composite group details.....	132
Show composite groups associated with a device group.....	133
Perform control operations on composite group devices	134
GNS repository	134
Automatic upgrade.....	134
Shared group definitions with GNS.....	135
GNS updates to group definitions	136
GNS setup.....	137

GNS in a multihost environment.....	138
GNS user authentication.....	138
Start the GNS daemon.....	138
Restart the GNS daemon.....	139
Manage the GNS daemon	139
Access groups in offline mode.....	139
View and release GNS daemon external locks.....	140
GNS daemon log file	140
GNS daemon options file	141
Remote mirror device group definitions	143
Identify GNS groups	144
GNS backup and restore mechanism	145
Manual backup to the GNS database.....	146
Manual restore from the GNS database.....	146
Troubleshoot GNS	146
List GNS data about groups	147
Show GNS data about groups.....	147
Chapter 7: Inline Compression.....	149
Inline compression overview.....	149
Inline compression control operations.....	149
Set or remove compression at storage group level.....	150
Remove compression at storage container resource level.....	151
Inline compression display and reporting.....	151
Report data compressibility.....	152
Efficiency reporting.....	153
Report array compression conversion status	155
Chapter 8: Fully Automated Storage Tiering.....	156
Fully Automated Storage Tiering	156
Storage Resource Pool management	156
Modify storage resource pools	157
Rename storage resource pools.....	158
FAST information reporting.....	158
List Service Level.....	159
Show specific Service Level	160
List storage groups.....	160
List Storage Resource Pools	161
Show specific Storage Resource Pool	164
Storage Resource Pool demand reporting	165
List Storage Resource Pools for thin devices	168
Chapter 9: Manage Configuration Changes.....	170
Configuration change management overview.....	170
Configuration change session rules	171
Symconfigure command execution formats and options	171
Verify and check sessions	172
Abort configuration session.....	173
Configuration change guidelines.....	174

Verify viable array configuration.....	174
Check for free physical disk space.....	174
Stop I/O activity on affected devices.....	174
Mixed array environments	175
Update host with device mapping information	175
Supported configuration operations	175
Set array attributes	176
View the array metrics.....	177
Set Service Level name.....	177
Rules and restrictions for set/reset Service Level name (HYPERMAX OS 5977 or higher)	177
Customize Service Level Response Time Multiplier.....	178
External disk group management.....	178
Create devices (HYPERMAX OS 5977 Q12016SR or higher)	178
Create devices (HYPERMAX OS 5977 lower than 5977 Q12016SR)	180
Restrictions when creating and managing devices	182
Expand device capacity.....	183
Copy devices	184
Copy a similar device.....	184
Convert devices	185
Device conversion restrictions.....	186
Valid thin device conversions	187
Convert a thin device.....	187
Thin device recommendations to maximize I/O activity	188
Set device attributes	188
Device attribute values.....	189
Set device identifiers.....	189
Identifier exclusions.....	190
View device identifiers	190
Delete devices.....	190
Device reservation management.....	191
Reserve devices.....	192
Commit changes on reserved devices	192
View reserved devices on an array.....	193
View details on reservation ID.....	193
Release reserved devices.....	193
Device pool management	194
Map CKD devices to CU image (HYPERMAX OS 5977 Q12016SR or higher).....	194
Unmap CKD devices from CU image (HYPERMAX OS 5977 Q12016SR or higher).....	195
Assign PAV alias addresses to CU image mapped devices	195
Remove PAV alias addresses from CU image.....	196
Map FBA devices to director ports.....	196
Obtain list of addresses.....	197
Unmap FBA devices from director port.....	197
I/O activity and unmapping devices.....	198
Managing PowerPath initiator and host registration.....	198
Introduce devices to a host	199
Introduce devices to Sun Solaris systems.....	199
Introduce devices to HP-UX systems.....	199
Introducing devices to IBM AIX systems	199
Introducing devices to HP Tru64 UNIX systems	200

Introducing devices to Windows systems	200
Set SRDF group attributes	200
Swap RA groups	200
Virtual Witness (vWitness).....	201
vWitness requirements.....	201
vWitness Management.....	202
Add director.....	202
Remove director.....	203
Port to director emulation support	203
Associate ports to director emulations.....	204
Disassociate ports from director emulations.....	205
Set port characteristics	206
SCSI protocol flags.....	208
Fibre protocol flags.....	209
Setting a port to show ACLX device	210
Report flag details.....	210
Chapter 10: Manage multiple iSCSI targets	212
Manage multiple iSCSI targets overview	212
Create IP interface on SE director port	213
Modify IP interface on SE director emulation	214
Delete IP interface on SE director emulation	215
Create iSCSI target on SE director emulation	216
Modify iSCSI target	218
Delete iSCSI target	220
Rename iSCSI target	221
Attach IP interface to iSCSI target	222
Detach IP interface from iSCSI target	223
Add IP route to SE director emulation	224
Remove IP route from SE director emulation	225
Example iSCSI configuration	226
Symconfigure command restrictions for iSCSI targets	226
Set online or offline state for iSCSI targets	227
Report iSCSI target information	227
List array IP addresses	227
List array iSCSI targets	228
List array IP routes	229
Expanded symcfg and symprd reporting for iSCSI targets	230
Expanded port group management for iSCSI targets	230
Report iSCSI target port information	231
Show port group with iSCSI targets	231
Show masking view with iSCSI targets	232
List HBAs with iSCSI targets (HYPERMAX OS 5977 or higher).....	233
List device information with iSCSI targets	233
List device assignments with iSCSI targets	233
List CHAP information with iSCSI targets	234
List login information with iSCSI targets	234
List storage group demand with iSCSI targets.....	235

Chapter 11: Manage storage environment for VMware VVols.....	237
Storage management for VVols overview.....	237
Solutions Enabler CLI support for VVol management	239
Create storage container for VVol management.....	239
Delete storage container for VVols.....	240
Modify description for storage container for VVols	240
Add storage resource to storage container for VVols.....	240
Modify storage resource for storage container for VVols.....	241
Remove storage resource from storage container for VVols.....	241
Create Protocol Endpoint (PE) devices for VVols.....	242
CLI command support for Protocol Endpoint (PE) devices for VVols	242
Report storage containers for VVols.....	243
Show storage container for vVols.....	243
List storage containers for VVols.....	244
Report PE and VVol devices.....	245
List PE devices and VVol devices.....	246
Unsupported operations/features for VASA protocol endpoints.....	247
Chapter 12: Array integration with RecoverPoint.....	248
RecoverPoint integration overview.....	248
RecoverPoint integration naming conventions.....	248
RecoverPoint device rules and restrictions.....	248
RecoverPoint integration control operations.....	250
RecoverPoint Environment setup - description and actions.....	250
RecoverPoint journal device creation - description and actions.....	252
RecoverPoint Environment expand - description and actions.....	253
RecoverPoint journal device deletion - description and actions.....	254
RecoverPoint device protection - description and actions.....	255
RecoverPoint device protection removal - description and actions.....	257
RecoverPoint repository creation - description and actions.....	258
RecoverPoint environment removal - description and actions.....	259
RecoverPoint integration reporting.....	260
List RecoverPoint clusters.....	260
List specific RecoverPoint cluster.....	261
List Recover Point protected storage group.....	263
Report RecoverPoint device types.....	264
Chapter 13: FAST.X.....	266
FAST.X	266
FAST.X overview	266
Solutions Enabler features supported with FAST.X.....	267
eDisks	268
eDisk encapsulation	268
Create external disk group	268
Delete an external disk group.....	269
Add an eDisk	269
Rules and restrictions for adding eDisks	270
Show external disk information	271

Remove an eDisk	271
List external spindle state.....	272
Verify eDisk state.....	273
Start drain on eDisk.....	273
Stop drain on eDisk	274
Activate eDisk.....	275
Chapter 14: Device masking with Auto-Provisioning Groups.....	277
Auto-provisioning groups.....	277
Provisioning limits.....	277
Create a masking view.....	278
Auto-provisioning session rollback.....	278
Create an auto-provisioning session.....	278
Discover host HBAs.....	279
List host HBAs.....	279
Create groups and views.....	279
Manage masking views.....	283
Delete masking views.....	283
Name groups and masking views	283
Back up and restore masking views.....	284
Copy groups and masking views.....	284
Manage storage groups.....	285
Add devices to storage group.....	285
Remove devices from storage group.....	286
Rename a storage group.....	286
Delete a storage group.....	287
List and show storage group information.....	287
Manage port groups.....	287
Add ports to port group.....	288
Remove ports from port group.....	288
Delete port groups.....	289
Copy a port group from array to array.....	289
List and show port group information.....	289
Fibre Channel ID lockdown.....	289
Locking down a Fibre Channel ID.....	290
Manage initiator groups.....	290
Add initiators to initiator group.....	290
Remove initiators from initiator group.....	291
Deleting initiator groups.....	292
Initiator group flags.....	292
Set HBA flags.....	293
Replacing a HBA.....	295
Rename a HBA.....	295
Using CHAP authentication.....	295
Verify the auto-provisioning database.....	297
Display auto-provisioning group information.....	297
Modified reporting commands for extended FA director port range.....	298
Display masking views.....	298
View auto-provisioning group details.....	300
View auto-provisioning device assignments.....	301

Display device list with no auto-provisioning assignments.....	302
List initiator group devices.....	302
View the HBA alias name.....	302

Part II: Querying and Reporting 304

Chapter 15: Configuration Query Operations..... 305

Configuration data overview.....	305
SCSI-level data	305
SCSI device list.....	306
Devices listed by array ID	306
List of devices without capacity	306
List of devices with WWN	307
Device list in pdevfile format	307
List HBA information	308
List device mapping information	309
List device identifiers.....	309
Array-level data.....	309
List arrays.....	310
List arrays by system resource demand.....	311
Data at Rest Encryption	312
View application registrations with array access	314
List host connections to arrays	315
List PowerPath host registration records.....	316
List host connections sorted by host names.....	317
Supported director configuration types	318
Array external locks	323
List all LRU cache management groups	325
Network configuration file	325
List memory board information	326
Mainframe CU image and split information reporting.....	326
List operating system patches.....	329
Environmental data	329
List device pools	331
View Snapshot policy details.....	332
Show thin pool rebalancing	335
List feature registrations and usage data	336
Device-level data	336
Device types.....	337
List array device information	337
External spindle devices	340
Report PowerPath device status.....	341
Report devices with non-native WWN.....	342
Change the device state.....	343
View device service states	344
Device emulation	344
List device encryption flags.....	344
List device emulation types.....	345
Show device details.....	345

Show clone state flags	346
Show disk geometry details	346
List DATA and SAVE devices	348
List thin device information	348
Verify device usage.....	353
Migrated devices	353
Filter list for device data	355
Device external locks	358
Disk-level data.....	358
List and show disk information.....	359
List SAS drives	359
List disk gaps	360
Reported actual disk capacity vs. rated disk capacity	360
List disk groups	360
List spare physical disks	362
List spindle information	363
External spindle information	365
Chapter 16: Events and Logs.....	366
Events and logs overview.....	366
Log option configuration	366
Array event monitoring using SYMCLI	369
Monitoring events.....	369
List events.....	369
Common audit log	370
Show audit log for single array.....	370
List audit log for specified time period.....	371
Custom audit log activity ID.....	372
Daemon log files.....	373
Event daemon	374
Chapter 17: XML Structured Output.....	375
XML structured output overview.....	375
XSLT: XML data transformations	375
Element-based XML	375
Set XML mode with SYMCLI	375
XML output using SYMCLI.....	376
Part III: Migrating Data.....	378
Chapter 18: Non-Disruptive Migration.....	379
Non-Disruptive Migration overview.....	379
Environmental requirements for Non-Disruptive Migration.....	380
Pre-migration rules and restrictions for Non-Disruptive Migration.....	380
Geometry Compatible Mode flag set with Non-Disruptive Migration.....	381
Non-Disruptive Migration for PowerMaxOS 5978 overview.....	382
Migration from a VMAX array.....	382
Migration from a VMAX3 or VMAX All Flash array.....	383
Environmental requirements for non-disruptive migration.....	384

Non-Disruptive Migration operational restrictions and state reference	385
Non-Disruptive Migration session states.....	385
Non-Disruptive Migration control actions and dependent migration states.....	386
Non-Disruptive Migration compatibility with other replication technologies.....	387
Non-Disruptive Migration restrictions with other SYMCLI commands.....	388
Non-Disruptive Migration control operations.....	389
Environment setup for Non-Disruptive Migration.....	389
Migration session creation for Non-Disruptive Migration.....	391
Cutover Non-Disruptive Migration session (only for Enginuity 5876 to HYPERMAX OS 5977, or to PowerMaxOS migration).....	394
Readytgt Non-Disruptive Migration session (only for HYPERMAX OS 5977 to PowerMaxOS migration).....	395
Commit Non-Disruptive Migration session.....	396
Cancel Non-Disruptive Migration session.....	397
Synchronizing devices for Non-Disruptive Migration session.....	398
Recover a failed Non-Disruptive Migration session.....	399
Remove environment for Non-Disruptive Migration session.....	401
List Non-Disruptive Migration session status	402
Non-Disruptive Migration examples	403
Example: Non-Disruptive Migration environment setup.....	403
Example: Typical Non-Disruptive Migration process (from Enginuity 5876 to HYPERMAX OS 5977, or to PowerMaxOS).....	404
Example: Typical Non-Disruptive Migration process (from HYPERMAX OS 5977 to PowerMaxOS 5978).....	406
Example: Suspending, restarting and recovering the Non-Disruptive Migration session.....	407
Example: Cancelling the Non-Disruptive Migration session.....	410
Example: Troubleshoot migration failure (device moved between storage groups).....	411
Example: Troubleshoot migration failure (masking view added).....	413
Chapter 19: Open Replicator	416
Open Replicator and arrays running HYPERMAX OS.....	416
ORS and host interaction	416
ORS rcopy concepts	417
ORS operational rules and limitations	418
ORS copying limitations	418
ORS device guidelines	418
ORS SAN setup requirements	419
ORS SYMCLI symsan support	419
Open Replicator session options	420
ORS hot pull copy session example	420
Open Replicator control operations	421
ORS and creating a device file	422
Create ORS copy session	424
Activate ORS session	426
Monitor ORS session status	428
Terminate an ORS session	434
Remove an ORS remote device from a session.....	435
Recreate an ORS session	435
Recover from failed ORS session	436
Restoring an ORS session	437

Open Replicator examples	438
Example: Perform an ORS hot pull operation	438
Example: Pull data online from IBM F20 to VMAX array	442
Example: Pushing online data from VMAX array to HDS 9960 array	448
Open Replicator operational restrictions and state reference	454
Device control operations allowed for ORS control devices	454
ORS operations allowed for replication session states.....	455
ORS operations allowed for device types.....	463
Part IV: Security related settings and procedures.....	464
Chapter 20: Log files and settings.....	465
Secure audit log.....	465
SYMAPI log files.....	465
Daemon log files.....	465
Chapter 21: Port Usage.....	466
Port usage.....	466
Configure Solutions Enabler server and event daemon ports.....	466
Chapter 22: Lockbox.....	468
Lockbox.....	468
Stable System Values (SSVs).....	468
Lockbox passwords.....	469
Password and SSV management.....	469
Index.....	470

1. Access Control Lists and Entries in the Access Control database.....	49
2. Example default configuration.....	66
3. Establish Administrator.....	68
4. Grant all permissions to non-pooled devices.....	68
5. Host-visible GNS state.....	135
6. GNS across multiple hosts.....	136
7. SRDF/Metro vWitness vApp and connections.....	201
8. VMAX3 VVol Architecture.....	237
9. High-level overview of FAST.X environment.....	266
10. Masking view overview.....	278
11. Masking view MV_1.....	282
12. Non-Disruptive Migration zoning	379
13. Configuration of a VMAX migration.....	382
14. Configuration of a VMAX3 or VMAX All Flash migration.....	383
15. Control array device pull operation.....	417
16. Control array hot pull using the symrcopy command.....	421

1. Typographical conventions used in this content.....	21
2. SYMCLI help options.....	24
3. Abbreviated command line actions.....	27
4. User authorization command file format.....	45
5. Initial access groups set.....	54
6. ACL setup for new array installation.....	54
7. Access Control Permissions: AccessType.....	61
8. Logical name formats.....	122
9. Location of daemon authorization file	138
10. Restarting the GNS daemon.....	139
11. Locations of GNS configuration options.....	141
12. Possible metric names (HYPERMAX OS 5977).....	176
13. Maximum device sizes	179
14. Maximum device sizes (HYPERMAX OS 5977).....	181
15. Thin device conversions.....	187
16. SCSI protocol port flags.....	208
17. Fibre protocol port flags	209
18. VVol architecture component management capability.....	239
19. Supported CLI commands for PE devices.....	242
20. Initiator group flags.....	292
21. Supported HBA flags.....	294
22. Pool Rebalancing Parameters.....	335
23. Log file configuration options.....	367
24. Options for managing daemon log files.....	373
25. Migration control actions and applicable migration states.....	386

26. Control and remote device guidelines	418
27. Device control operations allowed for ORS control device states with host.....	454
28. Allowable rcopy operations when the control device for PUSH session is in use as a TimeFinder/Mirror SOURCE.....	455
29. Allowable rcopy operations when the control device for PUSH session is in use as a TimeFinder/Mirror TARGET.....	456
30. Allowable rcopy operations when the control device for PULL session is in use as a TimeFinder/Mirror SOURCE.....	456
31. Allowable rcopy operations when the control device for PULL session is in use as a TimeFinder/Mirror TARGET.....	457
32. Allowable rcopy operations when the control device for PUSH session is in use as a TimeFinder/Clone SOURCE.....	458
33. Allowable rcopy operations when the control device for PUSH session is in use as a TimeFinder/Clone TARGET.....	458
34. Allowable rcopy operations when the control device for PULL session is in use as a TimeFinder/Clone TARGET.....	459
35. Allowable rcopy operations when the control device for PULL session is in use as a TimeFinder/Clone SOURCE.....	460
36. Allowable rcopy operations when the control device for PUSH session is in use as an SRDF R1 mirror.....	461
37. Allowable rcopy operations when the control device for PUSH session is in use as an SRDF R2 mirror.....	461
38. Allowable rcopy operations when the control device for PULL session is in use as an SRDF R1 mirror.....	462
39. Allowable rcopy operations when the control device for PULL session is in use as an SRDF R2 mirror.....	462
40. ORS operations allowed by device type.....	463

Please Read: Solutions Enabler and Enginuity 5876 document support

This guide documents the latest Solutions Enabler version management CLI capabilities for VMAX3, VMAX All Flash, and PowerMax arrays running HYPERMAX OS. This version of Solutions Enabler also supports arrays running Enginuity 5876, however management features specific to these arrays are not documented in this guide. These features, including Managing Quality of Service (*symqos*) are found in the Dell EMC Solutions Enabler Array Controls and Management CLI User Guide version 8.3 or higher.

References in this document to Enginuity version 5876 are limited to Chapter 19 "Non Disruptive Migration" where data migrations run between source arrays running Enginuity 5876 and target arrays running HYPERMAX OS or PowerMaxOS.

PREFACE

As part of an effort to improve its product lines, Dell EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your Dell EMC representative if a product does not function properly or does not function as described in this document.

NOTE: This document was accurate at publication time. New versions of this document might be released on Dell EMC Online Support (<https://support.emc.com>). Check to ensure that you are using the latest version of this document.

Purpose

This document describes how to use Dell EMC Solutions Enabler SYMCLI commands to configure and manage VMAX3 and VMAX All Flash arrays using HYPERMAX OS, and PowerMax arrays running PowerMaxOS.

- [Setting up and Configuring Storage Arrays](#) on page 23 describes how to use the SYMCLI configuration commands to set up the array devices.
- [Querying and Reporting](#) on page 304 describes how to use the SYMCLI query commands and various reporting tools to monitor and manage storage array performance.
- [Migrating Data](#) on page 378 describes how to use SYMCLI Non-Disruptive Migration and Open Replicator to migrate data between arrays.
- [Security related settings and procedures](#) on page 464 includes the procedures used to configure security parameters for Solutions Enabler deployment.

Audience

This document is part of the Solutions Enabler documentation set, and is intended for use by advanced command-line users and script programmers to manage various types of control operations on VMAX3, VMAX All Flash, and PowerMax arrays and devices using the Solutions Enabler SYMCLI commands.

Related documentation






The following documents provide additional Solutions Enabler information:

- | | |
|--|---|
| <i>Dell EMC Solutions Enabler, VSS Provider, and SMI-S Provider Release Notes</i> | Describes new features and any known limitations. |
| <i>Dell EMC Solutions Enabler Installation and Configuration Guide</i> | Provides host-specific installation instructions. |
| <i>Dell EMC Solutions Enabler CLI Reference Guide</i> | Documents the SYMCLI commands, daemons, error codes and option file parameters provided with the Solutions Enabler man pages. |
| <i>Dell EMC Solutions Enabler SRDF Family CLI User Guide</i> | Describes how to configure and manage SRDF environments using SYMCLI commands. |
| <i>Dell EMC Solutions Enabler TimeFinder SnapVX CLI User Guide</i> | Describes how to configure and manage TimeFinder SnapVX environments using SYMCLI commands. |

EMC Solutions Enabler TimeFinder Family (Mirror, Clone, Snap, VP Snap) CLI User Guide	Describes how to configure and manage TimeFinder Mirror, Clone, Snap and VP Snap environments using SYMCLI commands.
EMC Solutions Enabler SRM CLI User Guide	Provides Storage Resource Management (SRM) information related to various data objects and data handling facilities.
EMC VMAX All Flash and VMAX3 Family Security Configuration Guide	Describes how to securely deploy a VMAX3 Family (100K, 200K, 400K) or VMAX All Flash (250F, 450F, 850F, 950F) array with HYPERMAX OS.
VMAX Management Software Events and Alerts Guide	Documents the SYMAPI daemon messages, asynchronous errors and message events, and SYMCLI return codes.

Special notice conventions used in this document

Dell EMC uses the following conventions for special notices:

-  **NOTE:** Indicates a hazardous situation which, if not avoided, will result in death or serious injury.
-  **WARNING:** Indicates a hazardous situation which, if not avoided, could result in death or serious injury.
-  **CAUTION:** Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.
-  **NOTE:** Addresses practices not related to personal injury.
-  **NOTE:** Presents information that is important, but not hazard-related.

Typographical conventions


Dell EMC uses the following type style conventions in this document:

Table 1. Typographical conventions used in this content

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Used for full titles of publications referenced in text
Monospace	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Used for variables
Monospace bold	Used for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means "or"
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

Dell EMC support, product, and licensing information can be obtained as follows:

Product information	Dell EMC technical support, documentation, release notes, software updates, or information about Dell EMC products can be obtained at https://support.emc.com (registration required) or https://www.dellemc.com/en-us/documentation/vmax-all-flash-family.htm .
Technical support	<p>Dell EMC offers a variety of support options.</p> <ul style="list-style-type: none">• Support by Product — Dell EMC offers consolidated, product-specific information on the Web through the Dell EMC Online Support site. The Support by Product web pages (https://support.EMC.com/products) offer quick links to Documentation, White Papers, Advisories (such as frequently used Knowledgebase articles), and Downloads, as well as more dynamic content, such as presentations, discussion, relevant Customer Support Forum entries, and a link to Dell EMC Live Chat.• Dell EMC Live Chat — Open a Chat or instant message session with an Dell EMC Support Engineer.
eLicensing support	<p>To activate your entitlements and obtain your VMAX license files, visit the Service Center on Dell EMC Online Support (https://support.EMC.com), as directed on your License Authorization Code (LAC) letter emailed to you.</p> <ul style="list-style-type: none">• For help with missing or incorrect entitlements after activation (that is, expected functionality remains unavailable because it is not licensed), contact your Dell EMC Account Representative or Authorized Reseller.• For help with any errors applying license files through Solutions Enabler, contact the Dell EMC Customer Support Center.• If you are missing a LAC letter, or require further instructions on activating your licenses through the Online Support site, contact Dell EMC's worldwide Licensing team at licensing@emc.com or call:<ul style="list-style-type: none">○ North America, Latin America, APJK, Australia, New Zealand: SVC4EMC (800-782-4362) and follow the voice prompts.○ EMEA: +353 (0) 21 4879862 and follow the voice prompts.
Solve Online and Solve Desktop	<p>Solve provides links to customer service documentation and procedures for common tasks. Visit https://solveonline.emc.com/solve/products, or download the Solve Desktop tool from https://support.emc.com/ and search for Solve Desktop. From Solve Online or Solve Desktop, load the <i>PowerMax and VMAX</i> procedure generator.</p> <p> NOTE: You need to authenticate (authorize) the Solve Desktop tool. After it is installed, familiarize yourself with the information under Help.</p>

Your comments

Your suggestions help us improve the accuracy, organization, and overall quality of the documentation. Send your comments and feedback to: VMAXContentFeedback@emc.com

Setting up and Configuring Storage Arrays

This section describes how to set up and configure storage array devices using the Solutions Enabler SYMCLI commands.

Chapters include:

Topics:

- [Introduction](#)
- [Discovery](#)
- [User and Group Authorization](#)
- [Host-based Access Control](#)
- [Thin Device Management](#)
- [Grouping Devices](#)
- [Inline Compression](#)
- [Fully Automated Storage Tiering](#)
- [Manage Configuration Changes](#)
- [Manage multiple iSCSI targets](#)
- [Manage storage environment for VMware VVols](#)
- [Array integration with RecoverPoint](#)
- [FAST.X](#)
- [Device masking with Auto-Provisioning Groups](#)

Introduction

This chapter introduces the Solutions Enabler Symmetrix command line interface (SYMCLI), which is used to manage your storage environment.

Topics:

- [Solutions Enabler overview](#)

Solutions Enabler overview

The Dell EMC Solutions Enabler SYMCLI provides management capabilities for storage environments. The SYMCLI is a specialized library consisting of commands run at the command line, or used within scripts. These commands monitor device configuration and status, and perform control operations on devices and data objects within a VMAX-based storage environment.

SYMCLI commands are run from the host operating system command line (shell). The SYMCLI commands are built on top of SYMAPI library functions, which use system calls that generate low-level I/O SCSI commands to the storage arrays. To reduce the number of inquiries from the host to the storage arrays, configuration and status information is maintained in a host database file (called the Symmetrix configuration database; `symapi_db.bin` by default).

i NOTE:

TimeFinder and SRDF CLI operational commands are not discussed in this guide. Refer to the EMC Solutions Enabler TimeFinder Family (Mirror, Clone, Snap, VP Snap) CLI User Guide, Dell EMC Solutions Enabler TimeFinder SnapVX CLI User Guide, or the Dell EMC Solutions Enabler SRDF Family CLI User Guide for SYMCLI commands for replication operations.

i NOTE: When using PowerShell, any part of the SYMCLI command that requires a comma must be enclosed in quotes otherwise the command returns an error. See the following example:

SYMCLI:

```
symrdf addgrp -label Metro53 -rdfg 53 -sid 130 -dir 1G:5,2G:5 -remote_rdfg 53 -remote_sid 176
-remote_dir 1G:25,2G:25 -gige -nop
```

PowerShell:

```
symrdf addgrp -label Metro53 -rdfg 53 -sid 130 -dir "1G:5,2G:5" -remote_rdfg 53 -remote_sid
176 -remote_dir "1G:25,2G:25" -gige -nop
```

SYMCLI help options

SYMCLI provides the following multiple options to obtain top-level help using the command set:

Table 2. SYMCLI help options

Command	Description
symcli	Provides the version number of the installed command line interface.
symcli -h	Describes how to use the SYMCLI command.
symcli -v	Displays the SYMCLI commands and a brief description of each command.

Table 2. SYMCLI help options (continued)

Command	Description
<code>symcli -env</code>	Displays a list of the environment variables that can be set for a SYMCLI session.
<code>symcli -def</code>	Displays a list of the environment variables that are set for the current SYMCLI session.

Individual command help

Syntax

To display command line help for a specific command, use the following syntax:

```
command -h
```

Examples

To display help for the **symcfg** command, enter:

```
symcfg -h
```

Each command has its own man page for command line reference. To view the man page for the **symcfg** command, enter:

```
man symcfg
```

In a UNIX environment, the SYMCLI man page directory (`/usr/symcli/man/`) must be included in your MANPATH environment variable.

In a Windows environment, the default directory for man pages is `C:\Program Files\EMC\symcli\man`.

Device and object references

There are a number of different terms used with SYMCLI commands to identify and specify a VMAX array device:

- *PdevName* or `pd` — Indicates a physical (host) device name
- *SymDevName* or `dev` — Indicates array device name
- *LdevName* or `ld` — Indicates a logical device name
- `-devs` — Indicates a range or list of devices

Environment variables

Environment variables are set to streamline the command line session. Setting environment variables to commonly used values, eliminates the need to specify those arguments on the command line.

Display environment variables

Examples

To view a list of environment variables that can be set for the SYMCLI session, enter:

```
symcli -env
```

To view the current environment variables, enter:

```
symcli -def
```

Enable environment variables

Syntax

From UNIX host, set an environment variable, use the following syntax:

```
setenv variable_name value
```

From Windows host, set an environment variable, use the following syntax:

```
set <environment variable> = <value>
```

Examples

For UNIX host, to enable the `SYMCLI_VERBOSE` variable to set the default command behavior to always display command output with details:

 **NOTE: For UNIX host, `setenv` commands can be run in Cshell (`csh`).**

```
setenv SYMCLI_VERBOSE 1
```

For Windows host, to enable the `SYMCLI_VERBOSE` variable to set the default command behavior to always display command output with details:

```
set SYMCLI_VERBOSE=1
```

Disable environment variables

Syntax

From UNIX host, turn off an environment variable, use the following syntax:

```
unsetenv variable_name
```

From Windows host, turn off an environment variable, use the following syntax:

```
unset <environment variable>
```

Examples

From UNIX host, to turn off the verbose mode, enter:

```
unsetenv SYMCLI_VERBOSE
```

From Windows host, to turn off the verbose mode, enter:

```
unset SYMCLI_VERBOSE=
```

Display full group names

Description

By default, if group names are longer than the allowed space allotment, the name truncates in the output for the `symcg`, `symdg`, and `symsg list` commands.

Syntax

To display the full group name in the list command output, use the following syntax:

```
setenv SYMCLI_FULL_NAME sid
```

Example

To display full group name for array 567, enter:

```
setenv SYMCLI_FULL_NAME 100200000567
```

Preset names and IDs

Description

To reduce repeated key strokes for a group of commands that require the same argument, set the SYMCLI_DG and SYMCLI_SID to a preset value.

Syntax

To preset the device group name, use the following syntax:

```
setenv SYMCLI_DG dgName
```

To preset the array ID, use the following syntax:

```
setenv SYMCLI_SID array_id
```

Examples

To preset the array ID 100200000567 for a series of consecutive commands, enter:

```
setenv SYMCLI_SID 100200000567
```

To preset the device name for all -g arguments, enter:

```
setenv SYMCLI_DG MyDG
```

Command input and output tips

There are a number of different terms and abbreviations used with SYMCLI commands to identify and specify a VMAX array device:

Command line time-saving tips

Actions can be abbreviated, as follows:

Table 3. Abbreviated command line actions

Command	Abbreviated command
symcfg discover	symcfg dis
symcfg -sid 000002304324 sync	symcfg -sid 24 sync
symcfg LIST	symcfg list
sympd show /dev/rdisk/c2t1d1s2	sympd show c2t1d1s2

Truncated output display

In some cases, the output data from SYMCLI list commands exceeds the available column width, and returned data is truncated and appended with an asterisk (*). To view the full data, use the **show** command or verbose option (**-v**), when available.

Version compatibility information and restrictions

Compatibility between Solutions Enabler versions, can be set for output displays, running older version scripts, and client/server interoperability.

Set compatibility mode for older Solutions Enabler versions

Description

Compatibility mode is set globally using the `setenv SYMCLI_MODE` environment variable, or at the command line using the `-mode` option on any command. This specifies the command output reporting style to be compatible with prior SYMCLI versions. Possible values are from V80, V81, V82, V83, and V84.

Examples

For example, to set the compatibility mode to Solutions Enabler V8.0, enter:

```
symcfg list -mode V80
```

Attempts to specify a compatibility mode earlier than V8.0 fails with an Invalid mode error.

```
symcfg list -mode V76
'V76': Invalid mode specified.
```

Prior version Solutions Enabler scripts

Specialized scripts developed using a prior version of Solutions Enabler can be run in compatibility mode. Compatibility mode specifies the command output reporting style to be compatible with prior SYMCLI versions.



NOTE:

While running older scripts in compatibility mode, output that is only available in later releases may not be returned.

Client/server version compatibility

The latest Solutions Enabler version has additional restrictions limiting the interoperability between client and server with different versions of Solutions Enabler. The following restrictions are enforced in this release:

- SYMAPI client connections to servers running Solutions Enabler versions lower than V8.4 are rejected.
- SYMAPI server connections from clients running Solutions Enabler earlier than V8.0 are rejected with the error code:

```
SYMAPI_C_NET_VERSION_TOO_OLD
```

The remote connection is refused, as the client SYMAPI version is not supported by the server.

- SYMAPI server connections from clients with a Solutions Enabler version later than server are rejected with error code:

```
SYMAPI_C_NET_VERSION_TOO_NEW
```

The remote connection is refused, as the client cannot have a newer SYMAPI version than the server.

- Attempts to open a database file written by a version prior to V8.0 fail with the following error code:

```
SYMAPI_C_DB_VERSION_TOO_OLD
The SYMAPI database file cannot be used because
it was written by a version of SYMAPI that is no longer supported.
```

- When comparing client and server versions, only the major and minor versions are compared (for example: 8.0, 8.1). The edit and patch levels are ignored.

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^{10}	1,024	10^3	1,000
megabyte	MB	2^{20}	1,048,576	10^6	1,000,000
gigabyte	GB	2^{30}	1,073,741,824	10^9	1,000,000,000
terabyte	TB	2^{40}	1,099,511,627,776	10^{12}	1,000,000,000,000

Discovery

This chapter describes the SYMCLI discovery process.

Topics:

- [Discovery overview](#)
- [Discovery syntax options](#)
- [Verify host configuration database is synchronized](#)
- [Scan for new devices](#)
- [Connectivity authorization](#)
- [Configuring virtual disk mapping](#)
- [Display virtual disks as array devices.](#)
- [Host configuration database file](#)
- [Configuration data synchronization](#)

Discovery overview

The discovery process retrieves array, volume-level configuration, and status information. Discovery refers to the process where array, volume-level configuration, and status information is retrieved. Discovered configuration and status data for all arrays, as well as their directors and volumes, is maintained in a configuration database file on each host. Once the environment is discovered, information requests are made to retrieve array-level (high-level) data or device-level (low-level) information from it using SYMCLI commands.

Before using the SYMCLI commands, run the `symcfg discover` command to build the configuration (SYMAPI) database. Do after each Solutions Enabler or array operating system upgrade, and when presenting new devices to or removing old devices from the host.

Discovery syntax options

Description

During the first command line session, or if a configuration change has occurred, the configuration database must be built or rebuilt with the most complete and current information for all physical devices connected to the host. The `discover` command scans all SCSI buses, collects information about all the arrays and devices found, and rebuilds the database with the collected device information and parameters from all local and remotely attached devices. For more information about the configuration database, refer to [Host configuration database file](#) on page 33.

Syntax

To run a discovery operation, use the following syntax:

```
symcfg discover [-all | -pdev [-sid SymmID] | -sid SymmID] [-cache | -nocache]
```

Options

-pdev

Limits the discovery operation to only physical device information.

-sid

Limits the discovery operation to a specific array.

Examples

To scan the hardware and rebuild the database, enter:

```
symcfg discover
```

Verify host configuration database is synchronized

Description

The `symcfg verify` command verifies if the host configuration database file is synchronized with the configuration of an array. If they are in sync, the verify action returns code 0 (the `CLI_C_SUCCESS` value). If they are out of sync, the verify action returns code 24 (`CLI_C_NOT_IN_SYNC` value).

The `symcfg list -status` command lists all arrays connected to the host, and shows whether the configuration has changed.

Examples

To run the verify action for a specific array, enter:

```
symcfg verify -sid 814
```

To run the list action, enter:

```
symcfg list -status
```

Sample output

For the `symcfg verify` command:

```
The Symmetrix configuration and the database file are in sync.  
echo $status 0
```

For the `symcfg list -status` command:


```
                S Y M M E T R I X  
                Mcode  
SymmID      Attachment  Model   Version  Config Changed  Discovered  
000192600305 Local      VMAX    5977     Yes            Yes  
000192600321 Local      VMAX    5977     Yes            Yes  
000192600198 Remote     VMAX    5977     Yes            Yes  
000192600256 Remote     VMAX    5977     Yes            Yes  
000192600282 Remote     VMAX    5977     No             Yes  
000192600284 Remote     VMAX    5977     No             Yes
```

Scan for new devices

Description

When new devices are mapped to a host on nodes that have not yet been created, those nodes must exist before the discovery operation can find the devices along their path. The `scan` operation does the following:

- Searches the environment for devices accessible to a given host.
- Creates the nodes for the new or changed devices.
- Activates the necessary processing on the host system to recreate a list of accessible devices.

 **NOTE: The `symcfg scan` command is only supported on UNIX platforms.**

If changes are associated with array devices, follow the `scan` operation with the `discover` operation.

The `discover` operation scans all devices on the host looking for array devices and builds or rebuilds the array host database.

Examples

To scan for new devices, enter:

```
symcfg scan
```

To discover changes associated with array devices, enter:

```
symcfg discover
```

Connectivity authorization

Description

Some arrays may require authorization information to provide access to the array. Use the `symcfg authorization` command to supply this information for use in subsequent discovery operations. The `symcfg authorization` command lists, adds, updates, or deletes this connectivity information. The `update` action updates the password of an existing entry.

Syntax

To authorize connectivity, use the following syntax:

```
symcfg
  authorization list [-hyperv] [-v]
  authorization add | update -host HostName
    -username UserName -password PassWord
    [-namespace NameSpace] [-ahost HostName]
    [<-vmport | -port> port] [-vmware]
  authorization add | update -host HostName
    -username UserName -password PassWord [-hyperv]
  authorization delete -host HostName
    -username UserName [-namespace NameSpace] [-vmware]
```

Configuring virtual disk mapping

About this task

Solutions Enabler can resolve all virtual disks that have only one array device in the VMware datastore (supported with ESX 3.5 and higher). To use VMware virtual disk mapping, configure the host access credentials and copy the SSL certificate from the server to the VM OS.

Steps

1. Configure host credential using the following syntax:

```
symcfg auth add -host HostName -username root -password PassWord -namespace -vmware
```

2. Copy the SSL certificate. For example from the server: `/etc/vmware/ssl/rui.crt`) to the VM OS as: `/var/symapi/config/viclient_cert.pem`

3. When virtual disk mapping is configured, use following syntax to return the virtual disk information:

```
./syminq
```

Results

Device		Product			Device	
Name	Type	Vendor	ID	Rev	Ser Num	Cap (KB)
/dev/sda		VMware	Virtual disk	1.0	N/A	20971520
/dev/sdc		VMware	Virtual disk	1.0	N/A	8388608
/dev/sdd		EMC	SYMMETRIX	5773	8600022000	92160
/dev/sde		VMware	Virtual disk	1.0	N/A	8388608
/dev/dm-0		VMware	Virtual disk	1.0	N/A	8388608

NOTE:

An expected capacity difference is visible between the output of `syminq` and `sympd list` commands. When creating the virtual disk, it can use all the space or less space than the array device; so the `syminq` output may show smaller capacity than what `sympd list` output shows. In some cases, `syminq` shows larger capacity than `sympd list`, because the VMware datastore (or virtual disk) can have multiple array devices.

Display virtual disks as array devices.

Example

To display virtual disks as array devices for array 266, enter:

```
sympd list -sid 266
```

Sample output

```
Symmetrix ID: 000194900266
```

Device Name	Dir	Device
Physical	Sym SA :P Config	Attribute Sts Cap (MB)
/dev/sdl	00044 01E:0 TDEV	N/Grp'd ACLX RW 16
/dev/sdas	00168 07E:0 TDEV	N/Grp'd RW 28
/dev/sds	006DC 07E:0 TDEV	N/Grp'd NR 1934

Host configuration database file

The host configuration database (**.bin**) file, which is stored on the server host system, contains the physical configuration information of SCSI devices and array parameters that define your entire storage complex. More than one configuration database file may be required to support your operational needs.

The host configuration database is sometimes referred to SYMAPI database (because of how the file is named), or the Symmetrix database file. All of these names are referring to the same configuration database file, `symapi_db.bin`, described next.

Database file location

- On UNIX, the default pathname for the configuration database file is: `/var/symapi/db/symapi_db.bin`
- On Windows, the default configuration database path is: `C:\Program Files\EMC\SYMAPI\db\symapi_db.bin`
- On OpenVMS, the default configuration database path is: `SYMAPI$DB:symapi_db.bin`

An additional configuration database (`.bin`) file can be created to meet specific requirements.

NOTE:

In a multi-database environment, verify that the correct database is being used. Always confirm the current database before applying a command. For safe command line environments, it is recommended to use the common (default) database.

Database file lock

Solutions Enabler utilizes a configuration database locking file. The lock file is created automatically and is given the same name as the configuration database file, but it is appended with an `_xlock` suffix. For example, `symapi_db.bin_xlock`.

Solutions Enabler uses this lock file to serialize access to the configuration database file. It contains no data and is only used as a lock.

If you protect the `symapi_db.bin` file to restrict the users permitted to perform Solutions Enabler management operations, this lock file should be protected as well. Both `symapi_db.bin` and `symapi_db.bin_xlock` should be given the same level of protection.

Changing the current database file name

About this task

Steps

1. Check which host configuration database file is being used:

```
symcli -def
```

2. Modify the environment variable `SYMCLI_DB_FILE`.

For example, to change the database file to `sybackup_db.bin` from a UNIX host in C shell (csh), enter:

```
setenv SYMCLI_DB_FILE /var/symapi/db/sybackup_db.bin
```

To perform the same operation on Windows, enter:

```
set SYMCLI_DB_FILE=C:\Program Files\EMC\Symapi\db\sybackup_db.bin
```

Database location in client/server mode

For security reasons, in client/server mode the configuration database file must be located within the default database directory on the server host:

- On UNIX, the default pathname for the configuration database file is: `/var/symapi/db`
- On Windows, the default configuration database path is: `C:\Program Files\EMC\Symapi\db`

Database access modes

The SYMCLI commands utilize different modes to access the host configuration database file:

read/write — Commands that control and/or modify database parameters, read the database file into memory, and provide simultaneous modification of both the in-memory database and the database file. During this access cycle, the database file is locked.

read/no write — Commands that list or show database parameters, read the database file into memory and allow modifications to the in-memory database. No modifications to the database file occur. During the access cycle, the database file is not locked.

Command modes: Online and offline

SYMCLI commands are run in online or offline mode. Commands that execute in online mode, such as configuration control operations, automatically attempt to gather the latest state and mode information from the arrays, and update the in-memory database and

configuration database file on the host. If a configuration change has occurred, commands that execute in online mode will attempt to discover the changed entity to retrieve and load any updated configuration information.

Commands that can execute in offline mode, such as `symcfg list`, retrieve data exclusively from the configuration database.

Database configuration information

Example

To view basic configuration information about the current database in use, enter:

```
symcfg -db
```

Sample output

```
symcfg list -db
Type of SYMAPI Database           : Full
Host Node Name which discovered the DB : api33
Host OS Type which discovered the DB  : LINUX

Version of SYMAPI Library which discovered the DB : 9.0.X.X (Edit Level: 2001)
Version of SYMAPI Library which wrote the DB      : 9.0.X.X (Edit Level: 2001)
Min Edit Level of SYMAPI Lib Req'd to Read the DB : 2500

Time Database Was Last Synchronized      : Mon Feb  9 14:33:20 2018
Time Any Device Group Was Last Modified   : N/A
```

NOTE:

The above output includes the minimum edit level of the SYMAPI library required to read the database. Any SYMAPI library with this version or higher and edit level or higher can read the current database.

Configuration data synchronization

A `sync` operation of the host configuration database with the storage environment configuration, performed once a discovery operation has occurred, is less performance-intensive than a full discover operation.

A `sync` operation interrogates just the known arrays (previously discovered) that are accessible from the host, and updates the configuration and status information in the configuration database file. In addition, this synchronization can be limited to a specific array, where a configuration change may have occurred.

Update configuration information (sync operation)

Description

The `sync` operation refreshes the array configuration database file with data from the arrays. It does not update statistical information. [#unique_46](#) provides information about obtaining array performance statistics. The `symcfg sync` command updates data for all known entities in the configuration database file, but does not scan for SCSI devices on the host. That means that any newly configured physical devices are not added to the database during the synchronization. Synchronizing data is an array-specific operation.

 **NOTE:** If presenting new devices to a host or removing devices that a host sees, use the `discover` command before running the `sync` command.

Syntax

To perform a `sync` operation, use the following syntax:

```
symcfg sync
```

Options

-sid SymmID	A specific array
-rdf	SRDF information
-bcv	BCV information
-local	Local array information
-dirsts	Director status information
-snap	TimeFinder/Snap information
-cfgmgr	Disk space and array configuration metrics
-rcopy	rcopy information
-env_data	Environmental data
-fast	Fully-Automated Storage Tiering (FAST) information
-tier	Array tier definition information
-sg	Storage group information
-vpdata	Thin Provisioning data
-masking	Refreshes (synchronizes) the configuration database file with masking information gathered from the array.

Inhibit Database synchronization

Description

To force some commands to operate in offline mode, use the environment variable `SYMCLI_OFFLINE`, which inhibits accessing the array to update the database.

Examples

For example, to globally force these commands to their offline mode (`-offline` option) from a UNIX host in C shell (csh), enter:

```
setenv SYMCLI_OFFLINE 1
```

It may be necessary to refresh the database if executing commands while offline mode is enabled or execute commands that normally run in the offline mode, such as the `symcfg list`. Most display commands can be run in the offline mode, which provides the fastest response, and does not require access to the array.

User and Group Authorization

This chapter describes how to set up user and group authorization using the SYMCLI `symauth` command.

Topics:

- [User and group authorization overview](#)
- [Authorization settings management](#)
- [Monitor authorization](#)

User and group authorization overview

User authorization restricts the management operations individual users or groups can perform on an array. User authorization and the host-based Access Control database are independent utilities, but can be used together for maximum security. Refer to [Host-based Access Control](#) on page 49 for more information on the Access Control database.

The `symauth` command maps a user or group to a specific role, and allows the user or group access to specific operations on either an entire array or on individual components within an array. Authorization is configured independently for each array.

For more information about the `symauth` command syntax, refer to the *Dell EMC Solutions Enabler SYMCLI Command Reference Guide*.

User roles

A role determines the operations a user can perform. Unlike host-based access control, a user is assigned a particular role for the entire array or specified components within the array. For each array, a user or group can be assigned up to four roles. Roles are predefined and cannot be changed.

The following roles are defined in Solutions Enabler:

- **None** — Has no rights.
- **Monitor** — Performs read-only operations on an array excluding the ability to read the audit log or Access Control definitions.
- **PerfMonitor** — Includes Monitor role permissions and grants additional privileges within the performance component of Unisphere for VMAX application to set up various alerts and update thresholds to monitor array performance.
- **StorageAdmin** has the following rights:
 - Can perform all management operations on an array or on individual components within an array in addition to all monitor operations. This is the only role that can be given access rights to specific components within an array and is limited to virtualization domain users. Refer to [User ID formats](#) on page 38 for user ID format to assign Storage Admin role for virtualization domain user.
 - Can modify the GNS group definitions and monitor all operations (even if only granted rights to one component), and also has application performance monitor privileges.
 - Can modify storage groups (add and delete devices in storage group).
- **SecurityAdmin** — Performs security operations (`symaudit`, `symacl`, `symauth`) on an array in addition to all monitor operations. Users or groups assigned the SecurityAdmin or Admin roles can create or delete component-specific authorization rules. The SecurityAdmin also has all Auditor rights.
- **LocalRep** — Performs local replication operations such as SnapVX on an array.
- **RemoteRep** — Performs remote replication operations such as SRDF on an array.
- **DeviceManage** — Performs control or configuration operations on devices.
- **Admin** — Performs all operations on an array, including security operations and monitor operations. The Admin also has StorageAdmin rights, SecurityAdmin rights, and application performance monitoring privileges.
- **Auditor** — Grants the ability to view, but not modify, security settings for an array (including reading the audit log, `symacl list`, and `symauth`) in addition to all monitor operations. This is the minimum role required to view the array audit log.

List roles supported by the array

Syntax

To list the roles supported by an array, enter:

```
symauth list -roles
```

Assign multiple user roles

Examples

The `symauth` command is enhanced to assign multiple roles (up to four) separated with a '+' character.

```
StorageAdmin+Auditor+Monitor
```

NOTE: If any other role is combined with a role of None, the None role is ignored.

The `remove` command removes roles only for users assigned multiple roles, as shown in this example:

```
symauth -sid <SymmID> commit <<!  
assign user H:mars\smith to role StorageAdmin+Auditor;  
assign group D:Eng\Sup to role SecurityAdmin+PerfMonitor;  
assign user H:mars\smith to role Monitor;  
remove group D:Eng\Sup from role PerfMonitor;  
remove user H:mars\smith from role Auditor;
```

User authorization rules

When a user or group attempts access to an array or one of its components, Solutions Enabler performs the following authorization checks:

1. Verifies if user has access to the entire array.
2. Checks if group has access to the entire array.
3. Verifies if user has access to the StorageGroup component within the array.
4. Checks if group has access to the StorageGroup component within the array.

The StorageGroup component allows an Role Based Authentication Control rule to apply to a specific Storage Group or set of Storage Groups, without providing access to all SGs. 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

The rights granted by each check are combined together to form an overall set of rights for the user or group. For example, if a user is assigned to the SecurityAdmin role and also belongs to finance group, that is assigned the Monitor role, then the user is granted both SecurityAdmin and Monitor rights.

User ID formats

Users are identified by the user ID format, *Type:Qualifier\UserName or GroupName*.

Where:

- *Type* — Specifies the type of security authority used to authenticate the user. Authentication types are listed below.
- *Qualifier* — Specifies the domain or host that the user is logged into.

- *Name* — Specifies the username or groupname relative to that authority. It cannot be greater than 32 characters, and spaces are allowed if delimited with quotes.

User authentication types

Options

L

Indicates a user or group authenticated by LDAP. 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 SMC server. For example:

C:Boston\Legal — Indicates that user group Legal logged in through SMC sever Boston.

D

Indicates a user authenticated by a Windows domain. The qualifier specifies the domain or realm name. For example:

D:sales\putman — User putman logged in through the Windows domain sales.

D:jupiter\Sales — Group Sales logged in through the Windows domain jupiter.

H

Indicates a user authenticated (by logging in) to some host. On Windows, this corresponds to logging into a local account on the host. The qualifier specifies the hostname. For example:

H:jupiter\mason — User mason logged in on host jupiter.

H:jupiter\Sales — Group Sales logged in on host jupiter.

V

Indicates a user or group authenticated by the SMC server. For example:

C:Boston\Legal — Indicates that user group Legal logged in through SMC sever Boston.

Examples

Only Virtualization domain users or groups are allowed the role of StorageAdmin, with access to specific array components such storage groups or thin data pools. To assign thin pool access to a Virtualization domain use, enter:

```
symauth -sid 1234 commit <<!
assign group V:Host932jx\joe to role StorageAdmin for component ThinPool ThinPool:Sales-2;!
```

Username and groupname formats

Description

Within role definitions, user IDs or group IDs are either fully qualified (as shown previously), partially qualified, or unqualified. When the hostname or domain portion of the *UserName* parameter or *GroupName* parameter is an asterisk, the asterisk is treated as a wildcard meaning any host or domain. Examples of this include: H:*\user, D:*\user, and *\user. In all other cases, the asterisk is treated as a regular character.

Examples

- D:ENG\Sales — Fully qualified name with a domain and groupname.
- D:*\jones — Partially qualified name that matches username jones within any domain.
- H:HOST\Eng — Fully qualified name with a hostname and groupname.
- H:*\jones — Partially qualified name that matches username jones within any host.
- jones — Unqualified username that matches any jones in any domain on any host.

Wildcards

When using a wildcard (such as the asterisk), a given user or group may be matched by more than one mapping in the database. When searching for a role, the authorization mechanism uses the closest ID match that it can find.

- If an exact match (for example, `D:sales\putman`) is found, it is used.
- If a partial match (for example, `D:*\putman`) is found, it is used.
- If an unqualified match (for example, `putman`) is found, it is used; otherwise the user role is None.

When using the asterisk to allow wildcarding of user or group names, the `symauth` command may fail with an invalid argument error if it is used to establish a rule (assign or reassign) with a user or group that is deemed invalid through API validity checks.

Authorization settings management

The `symauth` command manages user and group authorization. This command enables and disables user and group authorization, sets enforcement mode, lists defined users and groups with their roles, and assigns roles to users and groups.

Current username identification

Options

-username

Displays the current username.

Examples

To display all groups to which the user belongs, enter:

```
symauth show -username
```

Sample output

This output displays all groups to which the user belongs.

```
Your current username:      D:Corp\Joe
Your current groupname: D:Corp\Finance
                        D:Corp\Sales
                        H:HostName\PowerUsers    ...
```

Group names are sorted alphabetically.

Manage Host IDs

Description

For improved security in user authorization, `symauth` supports managing Host IDs (Access IDs). These are the same host Access IDs used by Symmetrix ACLs.

Using the `symauth` command, the following operations can be performed on Host IDs:

- Obtaining Host IDs,
- Associating and disassociating Host IDs,
- Viewing and listing Host IDs,
- Importing Host IDs from existing `symacl` backup.

Syntax

```
symauth -unique [-passphrase [<PassPhrase> | -file <PassFile>]
                [-force]
```



```
symauth [-sid <SID>] list -hosts

symauth [-sid <SID>] commit [-v | -noecho] [-noprompt]
      [-file <CommandFile> | 'redirect stdin']

symauth -sid <SID> commit -acl_import -file <AclBackupFile>
```

where:

-acl_import

Specifies that Host ID values should be loaded from an existing Symmetrix ACL backup file. From this backup, Access Groups and Access IDs are interpreted as RBAC Host Names and Host IDs. Any Host Names that are already defined within the RBAC database are ignored.

-file

Specifies an existing file to be processed.

1. *<CommandFile>*: A command file to be processed for changes to the RBAC database.
2. *<PassFile>*: A file holding the passphrase to be used when generating a Host ID value for some host.
3. *<AclBackupFile>*: An existing Symmetrix ACL backup file to use when importing Host ID values to the RBAC database.

-force

If a Host ID (or symacl Access ID) is already defined for the host in question, this flag is needed to overwrite or replace that value.

-passphrase

Specifies a passphrase to be used in the Generation of a (new) Host ID. An alternative to supplying a passphrase on the command line is to provide it in a file on disk (the -file option).

-unique

Returns a secure 24-digit Host ID for the host being executed on. If necessary, one is defined – using either a supplied passphrase or certain hardware characteristics of the host.

-hosts

Lists the hosts that have Host IDs assigned with them.

Obtain Host ID

Description

The `symauth -unique` command can be used to obtain Host IDs (AccessIDs). This is the same operation that can be performed using the `symacl -unique` command.

Example

To obtain the host ID for the controlling host, enter:

```
symauth -unique
```

List hosts with Host ID

Description

The `symauth list -hosts` command is used to display the hosts that have had Host IDs associated with them

Example

To see hosts with associated Host IDs on array 041, enter

```
symauth list -sid 041 -hosts
```

```
Symmetrix ID: 000197802041
```

```
Host Name                               Host Access ID
```

```
-----  
...  
Host1234                *567  
Mercury                 *83D  
Venus                   *A39  
...
```

Associate/disassociate Host IDs

Description

The `symauth commit` command can be used to associate and/or disassociate the generated Host IDs with Host names. This operation requires a command file with entries to specify the one or more of the following operations:

- Associate a Host ID with a host name using `assign host <HostName> to hostid <HostID>`.
- Disassociate a Host ID from a host name using `delete host <HostName>`.
- Rename a host name in an existing Host ID using `Rename host <HostName> to <NewHostName>`.

Example

To associate a Host ID with a host name, enter:

```
symauth -sid 086 commit -file <CommandFile>
```

Import Host IDs

Description

The `symauth commit -acl_import` command is used to import Host IDs from an existing `symacl` backup file.

Example

To import Host IDs from an existing backup, enter:

```
symauth -sid 086 commit -acl_import -file ExampleBackup.File
```

User authorization

User and group role mappings should be defined before enabling authorization. At a minimum, there must be one mapping for an individual to the Admin or Security Admin, as user authorization can only be enabled if there is an individual (as shown by `symauth show -username` command) mapped to a role of either Admin or SecurityAdmin roles. These roles provide the ability to perform authorization control operations. Up to four roles can be assigned to a user group.

Set user authorization

Description

Assign user or group roles using the `symauth` command with role mappings supplied through a command file, or, on a UNIX host, supplied by a redirection of STDIN. Each line of the command file must contain a valid username or groupname with supported roles (refer to [User roles](#) on page 37 for supported roles), and end with a semicolon. For further details on using the command file, refer to [User authorization command file usage](#) on page 45.

Syntax

To create role mappings using a command file, use the following command syntax:

```
symauth -sid SymmID commit -file PathName
```

Examples

To create the user-to-role and group-to-role mappings, by supplying a redirection of STDIN (on UNIX host):

```
symauth -sid 1234 commit <<!
assign user H:jupiter\laura to role SecurityAdmin+PerfMonitor;
assign user D:Eng\neil to role Monitor;
assign user D:Eng\dave to role StorageAdmin+Auditor;
assign user D:Eng\steve to role Admin;
assign user D:Eng\paul to role PerfMonitor;
assign group V:Host932jx\joe to role StorageAdmin for component ThinPool ThinPool:Sales-2;
assign group D:Finance to role Monitor;!
```

NOTE: When using the asterisk (*) to allow wildcarding of user or group names, the `symauth` command may fail with an invalid argument error, if it is used to establish a rule (assign or reassign) with a user or group that is deemed to be invalid through API validity checks.

Enable user authorization

Description

User and group authorization is disabled by default. Any user can make changes to the authorization control data, including creating and removing user and group role mappings. Once authorization is enabled for an array, only users or groups with the Admin and SecurityAdmin role can change authorization control data. In addition, only a Virtualization Domain user assigned the role of StorageAdmin of an entire array can create or delete component-specific authorization rules. Refer to [User ID formats](#) on page 38 for assigning component specific authorization.

Only a user or group granted the role of Admin or SecurityAdmin can enable or disable authorization for an entire array.

Examples

To enable user authorization on array 097, enter:

```
symauth -sid 097 enable
```

NOTE: Specifying an array ID is optional for `enable` action.

User authorization enforcement

Description

Once user authorization is enabled, failed access attempts are enforced in one of two ways. An authorization failure can either result in operation failure or simply log a warning. The level of security enforced, is be set using the `symauth set enforcement` command.

Syntax

To set the security enforcement level, use the following syntax:

```
symauth -sid <SymmID> set enforcement <advise | enforce>
```

Options

advise

This parameter allows the operation to proceed when user or group authorization is denied, but generates a warning message (less secure) to the Solutions Enabler log file. This is an effective mode for validating your user and group role mapping without preventing access to array functionality.

enforce

This parameter, which is the default setting, causes the operation to fail when user or group authorization is denied (more secure), and logs the authorization attempt with its associated User or Group ID to the common audit log, `symaudit` log, and the Solutions Enabler log file.

 **NOTE:** Enforcement can also be set using a command file as described in [User authorization command file usage](#) on page 45.

Set Secure Reads

Description

The secure reads policy allows users with SECURITY_VIEW permissions to view the full set of authorization rules. All other users can only view rules affecting them. Enabling/disabling secure reads can be set using the `symauth set secure_reads <enable | disable>` command.

Syntax

To set the secure reads policy, use the following syntax:

```
symauth -sid <SymmID> set secure_reads <enable | disable>
```

Modify user authorization role mappings

Description

Modify user or group roles using the `symauth` command with role mappings supplied through a command file, or, on a UNIX host, supplied by a redirection of STDIN. Each line of the command file must contain a valid username or groupname with supported roles (refer to [User roles](#) on page 37 for supported roles), and end with a semicolon. For further details on using the command file, refer to [User authorization command file usage](#) on page 45.

Syntax

To modify role mappings using a command file, use the following command syntax:

```
symauth -sid SymmID commit -file PathName
```

Options


-file

This is the fully qualified path of the file containing the user-to-role mappings.

Examples

To modify user or group role mappings supplying a redirection of STDIN (on Unix host), enter:

```
symauth -sid 1234 commit <<!
assign user H:jupiter\laura to role Monitor;
assign user D:Eng\neil to role Admin;
assign user lauren to role StorageAdmin for StorGrp SG10;
reassign group Sales to role Auditor;
delete group Marketing forThinPool ThinPool:Sales-1;
reassign user D:Eng\bob to role Monitor;
reassign user V:Host932jx\joe to role StorageAdmin for ThinPool ThinPool:Sales-4;
!
```

 **NOTE:** When using the asterisk to allow wildcarding of user or group names, the `symauth` command may fail with an invalid argument error, if it is used to establish a rule (assign or reassign) with a user or group that is deemed to be invalid through API validity checks.

User authorization command file usage

A command file can be used to set user authorization, so a sequence of operations are previewed and committed at once. A command file is a simple text file formatted as follows:

Table 4. User authorization command file format

Action	format
Add a user to roles	assign user <i>UserName</i> to role <i>RoleNames</i> ;
Remove roles from user	remove user <i>UserName</i> from role <i>RoleNames</i> ;
Reassign an existing group to roles:	reassign group <i>GroupName</i> to role <i>RoleNames</i> ;
Delete a user	delete user <i>UserName</i> ;
Set enforcement	set enforcement [advise enforce];

NOTE: Assign operation is additive and does not replace an existing role. A user or group is allowed up to four authorization roles. Remove operation will remove individual roles for users with multiple roles.

Naming rules

- The following naming rules apply to user names and group names:
 - *UserName* or *GroupName* is the name of a user or a group (32-character maximum).
 - Spaces are allowed if the name is quote delimited.
 - Examples of a valid *Username* include: "D:domain\joe", "H:host\joe", "domain\joe", and "joe". For "domain \joe" it is interpreted as "D:domain\joe". For "joe", this *Username* is allowed regardless of domain or host.
- The following rules apply to *RoleNames*:
 - *RoleNames* are the roles assigned to or removed from a user or group.
 - Current valid roles include: *Admin*, *SecurityAdmin*, *StorageAdmin*, *Auditor*, *PerfMonitor*, *Monitor*, and *None*.
 - *RoleNames* are case insensitive.
 - To assign multiple roles to a user or group, place a plus sign "+" between roles (for example, **assign user User1 to role StorageAdmin+Auditor**).

Preview and commit user authorization actions

Description

To apply authorization actions (command file entries) to the array perform these operations on the command file:

1. Preview — This verifies the syntax and contents of the file.
2. Commit - This verifies and commits the contents of the command file to the array.

Syntax

To preview the command file entries, use the following syntax:

```
symauth -sid SymmID preview -file PathName
```

To commit the command file entries, use the following syntax:

```
symauth -sid SymmID commit -file PathName
```

Backup the authorization database

Description

The backup operation saves the contents of the user and group authorization database from the array to the specified file. Use the following command syntax to back up your user authorization database, where *SymmID* is the 12-character ID that specifies the array and *BackupFile* is the fully qualified path and filename of the backup that will be created:

Syntax

To backup the database with a fully qualified path and filename for the backup file, use the following syntax:

```
symauth -sid SymmID backup -f BackupFile
```

Restore the authorization database

Description

The restore operation re-initializes the user and group authorization database on an array, from a previously generated backup file, and re-enables user authorization. The specified file must be created by an earlier backup operation and can be from the same or a different array. Use the following command syntax to restore a previously created backup file, where *SymmID* is the 12-character ID that specifies the array to restore the file and *BackupFile* is the fully qualified path and filename of an existing backup file:

Syntax

To restore the database with the fully qualified path and filename for an existing backup file, use the following syntax and specify:

```
symauth -sid SymmID commit -restore -f BackupFile [-noprompt]
```

The restored file must assign the current user a role of Admin or SecurityAdmin; if it does not, the final restore step, which re-enables user authorization will fail. If this occurs, assign a role of Admin or SecurityAdmin to the current user and manually enable user authorization. Otherwise, have a user with Admin or SecurityAdmin privileges re-enable user authorization on the array.

Monitor authorization

Description

Use the monitoring operations to view the current user authorization settings on a specific array.

Examples

To list the user authorization policies in effect on an array, enter:

```
symauth -sid 025 list
```

To list the supported roles, enter:

```
symauth list -roles
```

To list the component types, enter:

```
symauth list -components
```

List user roles, names, and components

Description

The `list` operation retrieves current user authorization information from the array. If current information cannot be retrieved, the operation fails.

Syntax

To list user authorization information, use the following syntax:

```
symauth [-sid SymmID] [-offline]  
list -users [-by_domain | -by_role | -by_user] [-current_user]
```

Options

-offline

Retrieves cached authorization data on disk and does not communicate with the specified array. If there is no cached data on disk, no data displays (appearing as if there is no authorization data on the array).

-v

Use this option to request a multi-line display or when the length of names or components exceeds the 79-character limit.

Examples

To list the mappings of roles to all user and group names and their applicable components for array 1234, enter:

```
symauth -sid 1234 list -users
```

Sample output

The output shows the list the mappings of roles to user and group names and their applicable components for array 1234.

```
S Y M M E T R I X   A U T H O R I Z A T I O N   U S E R S  
Symmetrix ID: 00000001234  
  
Role                User/Group name                Component                Flags  
-----  
Admin               User  Mike                       N/A                      .  
SecurityAdmin      User  Smith                       N/A                      .  
StorageAdmin       User  H:mars\Jones                   N/A                      .  
                   User  H:saturn\Sally                 N/A                      .  
                   Group D:Eng\QA                       N/A                      .  
LocalRep            User  Jack                           StorGrp:testSG          .  
DeviceManage       User  Jack                           StorGrp:test2SG         .  
Auditor             User  H:mars\Jones                   N/A                      .  
                   Group D:Eng\QA                       N/A                      .  
PerfMonitor        User  D:Eng\Joe                       N/A                      N
```

Legend for Flags:
(E) : N = Rule has no effect since a different rule grants greater rights.
: . = Rule is active.

 **NOTE:**

- User `H:mars\Jones` has *StorageAdmin* role and which grants access to array 1234. This authorization rule is active as noted by the period (.) in the `Flags` column.
- User `D:Eng\Joe` is granted authorization to performance monitoring on the array, but the flag displays `N` (NoEffect) because another rule supersedes this one. User `D:Eng\Joe` also belongs to Group `D:Eng\QA` which is assigned the *StorageAdmin* role for the array and already grants performance monitoring rights to its group members.

Host-based Access Control

This chapter describes how to set up and perform Access Control actions using the `symac1` command.

Topics:

- [Host-based access control overview](#)
- [Create or modify access control data](#)
- [Create and manage access groups](#)
- [Create and manage limited access to access pools](#)
- [Create and manage access control entries](#)
- [Obtain access control information](#)
- [Verify a locked access control session](#)
- [Release locked access control sessions](#)
- [Access control strategies](#)
- [Backup and restore the access control database](#)

Host-based access control overview

The `symac1` access control command sets up and limits access to array resources (access pools). The command sets up access control mechanisms and changes access control entries through the Access Control database.

For more information about the `symac1` command syntax, refer to the *Dell EMC Solutions Enabler CLI Command Reference*.

Access Control database

An array-based access control database contains all the mechanisms or information to limit access to array access pools.

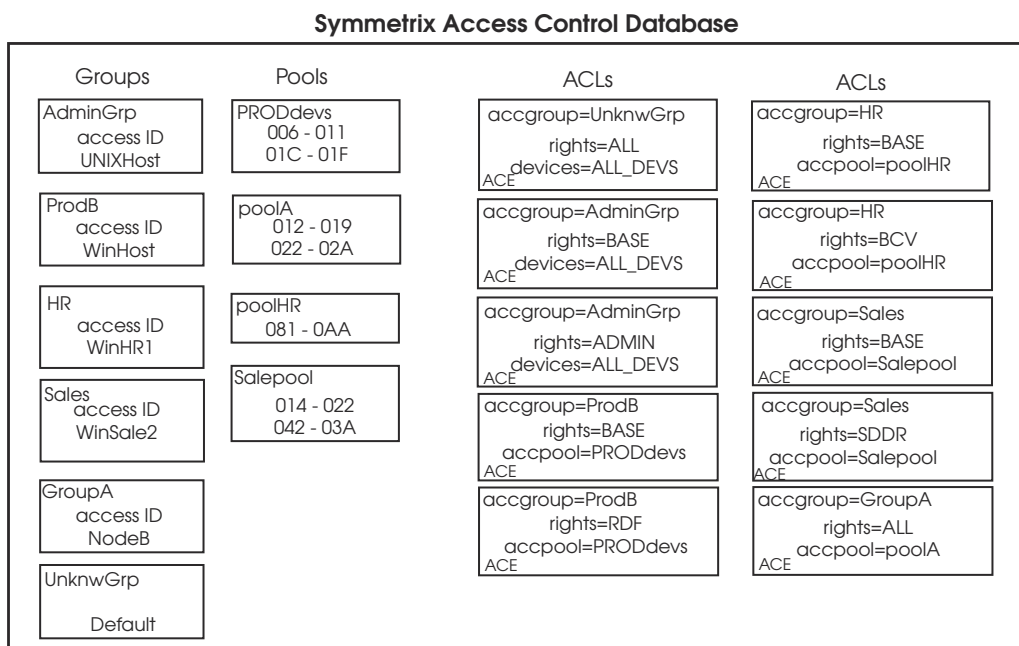


Figure 1. Access Control Lists and Entries in the Access Control database

The access control database contains the following access control components:

- **Access Control groups** — Unique *access IDs* and *names* are assigned (together) to hosts and then sorted into access control *groups* according to similar needs (determined by an Administrator). Access groups are allowed to act on access pools based on *permissions* (*access types*) granted by the Administrator. The unique host ID for open systems can be viewed by running the `symac1 -unique` command.
- **Access pools** — Permissions (or access types), such as BCV, SRDF, ADMIN, are assigned to allow a host to perform certain Solutions Enabler functionality on a specified set of devices. These sets of devices are referred to as access pools or *accpool*. Refer to [#unique_79/unique_79_Connect_42_SECTION_1920D6264A324C81AAED962AD74E227E](#) on page 61 for details on access types.
- **Access Control Entry (ACE)** — Once the group and access pool mechanisms are established, the ACEs are created, which grant permissions to these pools. The ACEs of the various access groups (along with groups and pools) are managed and stored in the Access Control database.
- **Access Control List (ACL)** — A group of ACEs that are associated with the same Access Control group.

Host access IDs

Symmetrix Access Control identifies individual management hosts using access IDs which are stored in a Lockbox. The Lockbox is associated with a particular host which prevents copying the Lockbox from one host to another. There are two different methods to generate the access IDs:

- Alternate access ID: The host's access ID can be generated at random or from a user-defined passphrase, then stored in a secure location on the local disk.

Alternate access IDs are supported for all platforms. See [Alternate access IDs](#) on page 50 for more information on alternate access IDs.

NOTE: It is strongly recommended that you use alternate access IDs on platforms where the hardware-based access ID is derived from a network interface MAC address.

- Hardware-based access ID (default): The host's access ID is derived from hardware characteristics of that host:
 - On x86_64 (64-bit Intel/AMD), and IA 64 platforms, a network interface MAC address is used.
 - On other platforms, characteristics of the host, such as a processor identifier, are used.

NOTE: When MAC addresses generate access IDs, the IDs may be unreliable or ineffective under some circumstances, including clustering environments, virtual environments, or following a hardware change. For added security on x86_64 (64-bit), IA64, and BS2000 hardware platforms, it is strongly recommended that you use alternate access IDs instead of hardware-based access IDs.

Alternate access IDs

Alternate access IDs are available for all platforms. When alternate access IDs are enabled, Solutions Enabler can:

- Randomly generate an access ID.
- Generate an access ID based on a user-chosen passphrase, where the passphrase is either:
 - Entered on the command line in an option.
 - Entered in a file, whose name is specified in the command line.

You enable alternate access IDs with the `SYMAPI_ALTERNATE_ACCESS_ID` option in the `<SYMAPI_HOME>/config/options` file.

Solutions Enabler securely stores the alternate access ID on the local disk in the Lockbox file. The `symac1` man page provides more information on the `symac1 -unique` command.

NOTE: Solutions Enabler access control changes must be made from an administrative host with ADMIN rights to the array and rights to make `symac1` changes.

If you only have one such administrative host, and you change its alternate access ID, once that change is made, the host can no longer make access control changes because the new access ID is not yet in an access group.

It is recommended that you enable a second administrative host prior to attempting to change a host's alternate access ID.

Enabling alternate access IDs

Steps

1. Add the following option to the `<SYMAPI_HOME>/config/options` file:

```
SYMAPI_ALTERNATE_ACCESS_ID = ENABLE
```

2. Run the `symacl -unique` command.

Solutions Enabler:

- Recognizes that the above option is set
- Generates an access ID if one does not already exist for the host
- Securely stores the access ID in the lockbox
- Displays the access ID

NOTE: If you run the `symacl -unique` command after enabling the options file setting, the new alternate access ID is different from the hardware-based access ID generated prior to enabling this option.

Any hardware-based access ID previously used to identify this host in an access group must be updated with the new alternate access ID using Solutions Enabler.

3. Use the `symacl` command to add the access ID to the appropriate access groups.

When an access ID is required on this host, the alternate access ID that was stored to disk is used.

Enabling alternate access IDs using a passphrase

Steps

1. Add the following option in the `<SYMAPI_HOME>/config/options` file:

```
SYMAPI_ALTERNATE_ACCESS_ID = ENABLE
```

2. Run the `symacl -unique` command using the `-passphrase` option.

The syntax when using the `-passphrase` option is as follows:

```
symacl -unique [-passphrase [Passphrase|-file PassFile] ]
```

NOTE: Passphrases can be 4 to 1000 characters long.

The following example shows how to activate an alternate access ID using a passphrase:

```
symacl -unique -passphrase Passphrase
```

The next example shows how to activate an alternate access ID using a passphrase stored in a file on the local disk:

```
symacl -unique -passphrase -file pathname
```

NOTE: In client/server mode, *pathname* names the file on the client host.

If no access ID already exists for the host, Solutions Enabler generates an access ID using the passphrase, securely stores it on the local disk, and displays it.

3. Use the `symacl` command to add the access ID to the appropriate access group.

When the access ID is required on this host, the alternate access ID that was stored to the disk is used.

Forcible generation of alternate access ID

Description

If the `SYMAPI_CLIENT_SIDE_ACCESS_ID` is enabled, then the `SYMAPI_ALTERNATE_ACCESS_ID` must be enabled. However, if the Alternate Access ID option is enabled on a host for the first time, along with the `SYMAPI_CLIENT_SIDE_ACCESS_ID` option and the client/server is set up, any command executed displays the error (Unable to obtain unique ID for host) because the client is attempting to send the client access ID to the server, but there is no client access ID to send,

Using the `symacl -unique -force` command in client/server mode to generate a client access ID is not permitted. To work around this, exit client/server mode on the client and execute the commands to generate a new access ID for the client.

Syntax

To forcibly generate a new alternate access ID using a supplied passphrase, use the following syntax:

```
symacl -unique -passphrase passphrase -force
```

To forcibly generate a new alternate access ID using a file that contains the passphrase, use the following syntax:

```
symacl -unique -passphrase -file PathnameToFile -force
```

NOTE: In the client/server mode, the *PathnameToFile* is on the client host, although the alternate access ID is activated on the server.

Disabling an alternate access ID

About this task

There are two methods to disable SYMAPI_ALTERNATE_ACCESS_ID option in the options file:

Steps

1. Change the following setting in the options file to: SYMAPI_ALTERNATE_ACCESS_ID = DISABLE or remove the line from the options file.
2. Run the **symacl -unique** command.

Results

This command recognizes that the option was reset, and disables the alternate access ID stored in the lockbox. The alternate access ID is retained in the lockbox in case you want to re-enable its use in the future.

Changing a host's alternate access ID

About this task

NOTE: Solution Enabler access control changes must be made from an administrative host with ADMIN rights to the array and rights to make symacl changes.

If you only have one such administrative host, and you change its alternate access ID, once that change is made, the host can no longer make access control changes because the new access ID is not yet in an access group.

Dell EMC recommends that you enable a second administrative host prior to attempting to change a host's alternate access ID.

For example, to change the access ID for Host-1:

Steps

1. Log in to another administrative host, such as Host-2.
2. Remove any existing Host-1 definitions from the access group for all arrays to which Host-1 has access.
3. From Host-1, follow the steps outlined in [Enabling alternate access IDs](#) on page 50 or [Disabling an alternate access ID](#) on page 52 to enable or disable the alternate access ID mechanism and obtain a new access ID.
4. From Host-2, add Host-1 back into its access group using its new access ID to any arrays to which it requires access.

Enabling hardware-based access IDs

About this task

NOTE: On IBM z/OS platforms, you must use job #14MSACL in the RIMLIB to generate and display the unique ID. #14MSACL takes the place of the symacl -unique command on this platform.

Steps

1. Confirm that the SYMAPI_ALTERNATE_ACCESS_ID option is disabled.
In the `<SYMAPI_HOME>/config/options` file, verify that the option is set to DISABLE, as shown below:
SYMAPI_ALTERNATE_ACCESS_ID = DISABLE
2. If the option is not set to DISABLE, do one of the following:
 - Edit the entry to set it to DISABLE.

- Remove or comment out the statement.
3. Run the `symac1 -unique` command to generate and display an access ID.
 4. Use the `symac1` command to add the access ID to the appropriate access groups.

Setting host access ID option

On the server host, the `SYMAPI_USE_ACCESS_ID` option controls the source of the access ID used for the client/server sessions:

`SYMAPI_USE_ACCESS_ID = CLIENT | SERVER | ANY`

The behavior of this option is as follows:

- `CLIENT` - The access ID supplied by the client host is used. If the client did not provide an access ID, operations fail. This can occur if the client is not configured to send its access ID.
- `SERVER` (default) - The server always uses its own access ID and ignores the access ID, if any, provided by the clients.
- `ANY` - The server uses an access ID provided by a singular client. If one is not provided, the server uses its own access ID.

Enabling or disabling client host access ID option

On the client host, the `SYMAPI_ALTERNATE_ACCESS_ID` option must be enabled to use this alternate access IDs:

`SYMAPI_ALTERNATE_ACCESS_ID = ENABLE`

Additionally, you must set the following option to control whether the client can send its own access ID to the server:

`SYMAPI_CLIENT_SIDE_ACCESS_ID = ENABLE | DISABLE`

The behavior of this option is as follows:

- `ENABLE` - The client sends its access ID to the server in client/server mode.
- `DISABLE` (default) - The client does not send its access ID to the server in client/server mode.

NOTE: After you enable the `SYMAPI_ALTERNATE_ACCESS_ID` and `SYMAPI_CLIENT_SIDE_ACCESS_ID` options, you must run the `symac1 -unique` command on the client host to generate the access ID and store it in the lockbox on the client host. Client/server mode MUST be enabled to perform this task.

Create or modify access control data

The `symac1` command creates or changes the access information in the Access Control database by committing a command file to the database.

The command file format contains various command entries terminated with a semicolon (;). The command syntax is case insensitive, but any variable parameter entered must be case sensitive.

The `symac1` command performs the following operations (specified in the command file):

- Creates new access groups
- Adds and removes access IDs to access groups
- Moves an access ID from one group to another
- Creates new access pools
- Adds and removes devices to access pools
- Deletes access pools and access groups
- Adds ACEs to grant access
- Removes ACEs to deny access

NOTE: After Access Control changes are made, a discover operation is required (`symcfg discover`).


Applying access control operations

About this task

To safely apply any Access Control operations (command file entries) to the access control database, perform the following `symac1` operations on the command file:

Steps

1. `Preview`
Verifies the syntax and accuracy of the contents of the entries in the command file.
2. `Prepare`
Performs the preview checks, but also verifies the validity of the requested access control modifications against the current state of the access control database in the array.
3. `Commit`
Performs both the `Preview` and `Prepare` checks and then commits the contents of the command file to the Access Control database.

 **NOTE:** It is not mandatory to execute a `Preview` or `Prepare` action prior to a `commit`. However, these operations can ensure that the `commit` action is not rejected, or they can be used to debug the command file entries.

Minimum access control configuration

To initialize the Access Control Database for a new array installation, create a host-based administrator and an access pin. This is usually performed by a Dell EMC Engineer and is described in the following example:

Table 5. Initial access groups set

Access groups	ID names
<code>AdminGrp</code>	<code>SunHost1 ACCPIN</code>
<code>UnknwGrp</code>	<code>unknown</code>

The name `SunHost1` shown in the table (associated with `AdminGrp`) is the ID name that designates the machine where access controls are administered. To perform `prepare`, `commit`, and `release` actions, an access ID (PIN) with the name **ACCPIN** is created in the delivered setup. This ID name is also used for the value setting of `SYMCLI_ACCESS_PIN`. For more information, see [Create and manage access groups](#) on page 55.

Initial ACL setup

As shown in the following table for the delivered setup described in [Minimum access control configuration](#) on page 54, access group `AdminGrp` has access to all devices (`ALL_DEVS`) and this group is granted both `ADMIN` and `ALL` permissions. Also in the initial setup, a group called `UnknwGrp` is established with permissions to all devices in the array. For initial system usage, this gives all unknown hosts `BASE` permissions to all devices (`ALL_DEVS`), until it becomes clear what restrictions should be established.

The initial `UnknwGrp` setup also grants `ALL` permissions to any devices not already assigned to an access pool (i.e. `!INPOOLS`).

Table 6. ACL setup for new array installation

Access groups	permissions (access types)	Access pools
<code>AdminGrp</code>	<code>ADMIN</code> <code>ALL</code>	<code>ALL_DEVS</code>
<code>UnknwGrp</code>	<code>BASE</code>	<code>ALL_DEVS</code>
<code>UnknwGrp</code>	<code>ALL</code>	<code>!INPOOLS</code>
<code>EmcIntrnlGrp</code> - used for internal array operations only	<code>ALL</code>	<code>ALL_DEVS</code>

 **NOTE:** The access type `ALL` excludes `ADMIN` privileges.

VLOGIX permission behavior

During initial setup of the system, the access group `UnknwGrp` (with Access Type `ALL` for `NON-POOLED` Devices) is present, as shown in [Table 6. ACL setup for new array installation](#) on page 54. The `VLOGIX` privilege returns as true since access is granted access to all devices in the array. Initially, because no pools are present, the `VLOGIX` privilege is associated implicitly with all devices by this ACL.

Once you create an access pool and add a device to it, the `VLOGIX` privilege is no longer implicitly associated with all devices and a check for the `VLOGIX` privilege will now fail. For the `UnknwGrp` to still have `VLOGIX` privilege, that privilege must be explicitly granted to the `UnknwGrp` and associated with `ALL_DEVS`.

Create and manage access groups

Various sets of users tend to use the same applications that utilize common Solutions Enabler features from a given host. System users typically share the same device resources and require access permissions to these shared devices. For shared devices, hosts are registered in groups identified with a group name, which serves as a root for all ACEs in the group (see [Figure 1. Access Control Lists and Entries in the Access Control database](#) on page 49).

Access groups contain groups of access IDs and their ID names. Any access ID and name can only belong to one group and are paired together into the database. For ease of management, it is highly recommended to use 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 is used to create access control entries (ACEs).

Verify administrative authority

Syntax

Prior to making any access control changes, verify the host used for ACL setup has the administrative authority. To verify the host, use the following syntax:

```
symacl list -v [-sid SymmID|ALL]
```

Create an access group

Examples

To add access groups named `HR` and `Sales` in array `12345`, enter in the command file:

```
create accgroup HR;  
create accgroup Sales;
```

To commit access groups to the database using the file `addnewgroups.cmd`, enter:

```
symacl -sid 12345 commit -file addnewgroups.cmd
```

At this point, a prompt may display for a 4 - 12 character access PIN (if environment variable `SYMCLI_ACCESS_PIN` is not already set to this PIN value).

Add host access IDs to a group

Description

The host access `Id` is case sensitive, allows alphanumeric characters, underscores, dashes and no spaces. The ID name length allows 1 - 31 characters.

Syntax

To add a host access ID to an access group, use the following syntax in the command file:

```
add host accid Id name IdName to accgroup GroupName
```

NOTE: The host access ID is encrypted; so choose an *IdName* that is easy to remember.

Examples

To add `winHost` with an access ID `73900158-06174491-16225515` to the `ProdB` group, enter in the command file:

```
add host accid 73900158-06174491-16225515 name WinHost to accgroup ProdB;
```

NOTE: To preserve access ID security for any host, IDs are encrypted. Use the `symacl -unique` command to get the encrypted value.

To commit the file (`addnewgroups.cmd`) and add the host access IDs to the specified groups, enter:

```
symacl -sid 12345 commit -file addnewgroups.cmd
```

NOTE: To update the array configuration after Access Control changes are complete, run the `symcfg discover` command. changes

User access IDs (PINs) for AdminGrp

Description

A user access ID is a PIN that allows a host to perform commit, prepare, or release operations to an array as the `AdminGrp`. When a host attempts a commit, prepare, or release operation as the `AdminGrp`, the `symacl` command prompts for this PIN.

The user access `Id` is case sensitive, allows alphanumeric characters, underscores, dashes and no spaces. The ID name length allows 1 - 31 characters.

Syntax

To add the user access ID for the **AdminGrp** access group, use the following command syntax in the command file:

```
add user accid Id name IdName to accgroup AdminGrp
```

NOTE: The host access ID is encrypted; so choose an *IdName* that is easy to remember.

Examples

To add the administrative user access ID (`1234PIN`) for `JOEPIN`, enter in the command file:

```
add user accid 1234PIN name JOEPIN to accgroup AdminGrp;
```

NOTE: User access IDs are set only for the `AdminGrp` access group. Setting an ID for other groups will not return an error and may not prompt for a PIN.

To commit the file `addnewgroups.cmd` and add the user access IDs to the specified groups, enter:

```
symacl -sid 12345 commit -file addnewgroups.cmd
```

Edit and manage an access group

Once access groups are established, access IDs or ACEs can be removed from a group, IDs can be moved between groups, or groups can be deleted.

 **NOTE:** To update the array configuration after Access Control changes are complete, run the `symcfg discover` command. changes

Remove an access ID from a group

Syntax

To remove an access ID from an access group, use the following syntax in the command file:

```
remove accid name IdName from accgroup GroupName
```

Examples

To remove user `HRUser2` from the `HR` group, enter in the command file:

```
remove accid name HRUser2 from accgroup HR;
```

To commit the file (`removeaces.cmd`) and remove the specified user, enter:

```
symacl -sid 12345 commit -file removeaces.cmd
```

Move an access ID to another group

Syntax

To move an access ID from one access group to another, use the following command syntax in the command file:

```
move accid name IdName to accgroup GroupName
```

Examples

To move user `HRUser1` to `GroupA`, enter in the command file:

```
move accid name HRUser2 to accgroup GroupA;
```

To commit the file (`moveaces.cmd`) and move the specified user, enter:

```
symacl -sid 12345 commit -file moveaces.cmd
```

Remove all ACEs from a access group

Syntax

To remove all ACEs from an access group, use the following syntax in the command file:

```
remove aces from accgroup GroupName
```

Examples

To remove all ACEs from the `HR` group, enter in the command file:

```
remove aces from accgroup HR;
```

To commit the file (`moveallaces.cmd`) and remove all ACEs from the specified group, enter:

```
symacl -sid 12345 commit -file moveallaces.cmd
```

Delete a access group

Syntax

To delete an access group from the database, use the following syntax in the command file syntax:

```
delete accgroup GroupName;
```

Options

`remove_aces=true`

If ACEs have not been removed from the group, this option removes all the ACEs as part of the `delete` operation.

Examples

To delete access group `HR` and any ACEs in the group, enter in the command file:

```
delete accgroup HR remove_aces=true;
```

To commit the file (`deletegroup.cmd`) and remove the specified user, enter:

```
symacl -sid 12345 commit -file deletegroup.cmd
```

Create and manage limited access to access pools

Access pools are groups of devices controlled by access groups. When the various devices used by a host application must function as non-shareable resources, the target devices must be identified and assigned into an access pool for protection. Once an access pool is created, the pool can be a target to create access control entries (ACEs). More than one access group can access a pool with different permissions. For example, group `AdminGrp` might access `PoolA` with `ALL` permissions, while group `HR` could access the same `PoolA` with just `BASE` permissions.

NOTE:

Once an access pool is created, any host in the access group `UnknwGrp` is denied access to the `symmask` and `symaccess` commands. In addition, when a host in `UnknwGrp` calls the `symmaskdb list database` or `symaccess list view -detail` commands, only its own devices will be returned in the list. To provide hosts in access group `UnknwGrp` with full access to the `symmask` command, grant the access type `VLOGIX` (for `ALL DEVS`) to access group `UnknwGrp`. For more information on granting access types with full access to the `symmask` and `symaccess` commands, refer to [Create and manage access control entries](#) on page 60.

Verify administrative authority

Syntax

Prior to making any access control changes, verify the host used for ACL setup has the administrative authority. To verify the host, use the following syntax:

```
symacl list -v [-sid SymmID|ALL]
```

Create an access pool

Examples

To create an access pool named `PoolB` in array `12345`, enter in the command file:

```
create accpool PoolB;
```

To commit the file `addnewpool.cmd` and add a new pool, enter:

```
symacl -sid 12345 commit -file addnewpool.cmd
```

Add devices to access pool

Syntax

To add the specific devices to a pool, use the following command syntax in the command file:

```
add dev StartDevName[:EndDevName] to accpool PoolName
```

Examples

To assign devices 0A through 19 to `PoolB`, enter in the command file:

```
add dev 00A:019 to accpool PoolB;
```

To commit the file `addnewdevices.cmd` and devices to the pool, enter:

```
symacl -sid 12345 commit -file addnewdevs.cmd
```

 **NOTE:** To update the array configuration after Access Control changes are complete, run the `symcfg discover` command. changes

Edit and manage access pools

Once access pools are created, devices can be removed from a pool, or pools can be deleted.

 **NOTE:** To update the array configuration after Access Control changes are complete, run the `symcfg discover` command.

Remove devices from access pool

Syntax

To remove specific devices from a pool, use the following command syntax in the command file:

```
remove dev StartDevName[:EndDevName] from accpool PoolName
```

Examples

To remove devices 16 through 19 from `PoolB`, enter in the command file:

```
remove dev 016:019 from accpool PoolB;
```

To commit the file `removedevs.cmd` and remove the specified devices from the pool, enter:

```
symacl -sid 12345 commit -file removedevs.cmd
```

Delete access pool

Syntax

To delete an access pool from the database, use the following syntax in the command file syntax:

```
delete accpool PoolName;
```

Options

`remove_aces=true`

Removes any ACEs associated with the pool.

Examples

To delete access pool `PoolB` and any associated ACEs, enter in the command file:

```
delete accgroup PoolB remove_aces=true;
```

To commit the file `deletepool.cmd` and remove the specified user, enter:

```
symacl -sid 12345 commit -file deletepool.cmd
```

Create and manage access control entries

Syntax

Access Control Entries (ACEs) grant permissions to groups and pools. Once access groups and pools are created, ACEs are created. A group can have multiple permissions and pools, which requires an ACE for each permission. ACEs determine and grant, to the access group, the permissions (*AccessType*) for access to a specified pool or all devices not in a pool.

Verify administrative authority

Syntax

Prior to making any access control changes, verify the host used for ACL setup has the administrative authority. To verify the host, use the following syntax:

```
symacl list -v [-sid SymmID|ALL]
```

Grant permissions to access group

Syntax

For command the file, use the following syntax:

```
grant access=AccessType to accgroup GroupName for accpool PoolName | ALL | NON-POOLED devs
```

To commit permissions granting, use the following syntax:

```
symacl -sid SymID commit -file FileName
```

Options

AccessType

Specifies permissions to the SYMCLI or SYMAPI features, or functionality granted to selected devices

Available access control permissions

Table 7. Access Control Permissions: AccessType

Permissions (<i>AccessType</i>)	Description	SYMCLI commands affected
ADMIN ²	Grants administrator privilege to grant/deny access control entries to hosts and users.	symacl <prepare, commit, release, list, show>
ADMINRD ²	Grants read-only access to all access control information.	symacl <list, show>
ALL ²	All possible access types granted except ADMIN and ADMINRD. Must be directed to ALL devices.	All
BASE ⁴	Allows the discovery of devices and to obtain states and statistics from the array (directors and devices).	Base component commands <ul style="list-style-type: none"> • symevent, symcfg, symaudit, symdg, symdev, symcg <list, show>, symdisk, symipsec (Policy priority list, Policy get, and IPSec statistics get operations) • symmigrate query and list • symlmf query, list, and show
BASECTRL	Allows base control operations on devices and device groups.	symdg <controls>, symcg <controls>, symdev <controls>, symsg <controls>
BCV	Allows TimeFinder BCV and TF/Clone control and status operations.	symbcv, symmir <controls>, symclone <controls>
CACHCTRL	Allows Cache control operations concerning LRU partition management.	symqos <set LRU >
CFGDEV	Allows powerful configuration control operations that manage various types of configuration changes on devices in the array.	symconfigure Allows the following types of operations: <ul style="list-style-type: none"> • Convert device configurations (BCV, SRDF, DRV) • Convert device configuration (changing mirroring) • Set device attributes • Set device emulation • Metadevice management
CFGSYM ²	Allows access to set array attributes, set port flags, and swap RA groups with symconfigure command. It also affects the symmaskdb , symaccess , and symrdf commands as specified in the next column. Must be directed to ALL devices.	symconfigure, symmaskdb, symrdf, symaccess Allows the following types of operations: <ul style="list-style-type: none"> • Spare management • SAVE device pool create/delete • SAVE device pool member management • Enable/disable SRDFA • Set matrix across the array • Add/remove dynamic SRDF groups • Change director supporting a dynamic SRDF group • Set the link limbo value for a dynamic SRDF group

Table 7. Access Control Permissions: AccessType (continued)

Permissions (<i>AccessType</i>)	Description	SYMCLI commands affected
		<ul style="list-style-type: none"> • Modify User Authorization settings via symauth • Converting thin device (adding SRDF) • Setting thin dev attributes • Bind/unbind thin device (not supported by symconfigure) • Configuring thin metadevices • Allocate/free thin device • Creating thin device pools • Adding devices to thin pools • Enabling/disabling devices in thin pools • Removing devices from thin pools • Deleting thin pools • symipsec policy add, replace/modify, and delete • Initialize the VCMDB • Convert the database type • Restore the database • Manage storage groups • Manage port groups • Manage initiator groups • Manage masking views
CHECKSUM	Allows array device Double Checksum operations.	symchksum <controls>
CREATEDV ²	Allows the creation and deletion of array devices (part of symconfigure).	symconfigure <ul style="list-style-type: none"> • Create new devices • Configure disk space (create/map/mask/meta/attributes) • Delete devices • Create thin devices
DIRCTRL ²	Allows you to take directors and their ports offline and online. Must be directed to ALL devices. Also allows you to set CHAP authentication.	symcfg <online/offline> <ul style="list-style-type: none"> • Online/offline RA directors • Online/offline front-end ports symconfigure <ul style="list-style-type: none"> • Set port flags and attributes symconnect, symaccess <ul style="list-style-type: none"> • CHAP authentication
ECC ¹⁻²	Allows the ECC Symmetrix agent to run on the requested host.	Not applicable
OPTMZR	Allows user-configurable attributes that may affect the Optimizer and Symmetrix Migration behavior.	symoptmz <controls> symmmigrate <controls> symfast <controls>
POWRPATH ¹⁻²	Access to PowerPath-directed devices in an SRDF consistency group. Must be directed to ALL devices.	Not applicable

Table 7. Access Control Permissions: AccessType (continued)

Permissions (<i>AccessType</i>)	Description	SYMCLI commands affected
QOS	Allows the execution of Quality of Service (QOS) performance control operations to manage copy priorities. Excludes LRU cache control functionality.	symqos <set pace>
RCOPY	Manage Open Replicator sessions.	symrcopy
RDF	Allows SRDF control and set operations.	symrdf <control>
RPA	For RecoverPoint splitter to operate correctly, you are required to grant BASE, BASECTRL, RPA, and CFGDEV access types.	Not applicable
SDDF	Allows the DeltaMark (Change Tracker) functionality that monitors track changes.	symchg
SDR	Allows mapping/unmapping of devices to directors/ports for the Symmetrix Disk Reallocation (SDR) feature, Device Masking, and Auto-provisioning Groups.	symconfigure <ul style="list-style-type: none"> • Manage CKD aliases • Map/unmap devices • Map/unmap thin device symmask, symmaskdb, symaccess <ul style="list-style-type: none"> • mapping and unmapping
SNAP	TF/Snapshot allows the creation and management of virtual copy sessions between a source device and multiple virtual (VDEV) target devices. Allows for creation of snapvx sessions	symsnap symsnapvx
VLOGIX ¹⁻⁴	Enables access to Device Masking devices and Auto-provisioning Groups.	symmask, symmaskdb, symaccess

1. See the appropriate product documentation for use of these access types.
2. These access types must be granted to either ALL devices or all NON-POOLED devices in an array.
3. This access type must be granted to ALL devices in an array.
4. BASE access allows the creation/modification of symaccess group types, which are not part of a view. VLOGIX access is required to create or manipulate items within a view

Grant BASE permissions to a group

Examples

Grant BASE permissions to a group before granting permissions to any value-added, non-base feature. To grant BASE permissions to access group `ProdB` for all devices, enter in the command file:

```
grant access=BASE to accgroup ProdB for ALL devs
```

To grant BASE permissions to access group `ProdB` for `poolPRODdevs`, enter in the command file:

```
grant access=BASE to accgroup ProdB for accpool PRODdevs
```

To commit the file `grantbaserights.cmd` and grant BASE permissions, enter:

```
symacl -sid 12345 commit -file grantbaserights.cmd
```

i **NOTE:** When restricting a host to BASE access to a pool of devices, all devices mapped to the host are visible along with any devices that are not mapped to the host, but are in the pool. To configure a host to see only devices that are in the pool, map only those devices to the host. In addition, remote (SRDF) arrays and their devices are discovered, and the application registration database and the audit logs cannot be accessed since these may contain data relevant to other hosts.

Grant SRDF permissions to a group

Examples

To grant SRDF permissions to access group `ProdB` for pool `PRODdevs`, enter in the command file:

```
grant access=RDF to accgroup ProdB for accpool PRODdevs
```

To commit the file (`grantsrdrfights.cmd`) and grant SRDF permissions, enter:

```
symacl -sid 12345 commit -file grantsrdrfights.cmd
```

Grant BCV and SDR permissions to a group

Examples

To grant BCV and SDR permissions to access group `HR` for pool `poolHR`, enter in the command file:

```
grant access=BCV, SDR to accgroup HR for accpool poolHR;
```

To commit the file (`grantbcvsdrrights.cmd`) and grant BCV, and SDR permissions, enter:

```
symacl -sid 12345 commit -file grantbcvsdrrights.cmd
```

Grant BASE permissions to a non-registered host

Examples

To make any non-registered (*unknown*) host have BASE permissions to access group `UnknwGrp` for all devices in the array environment, enter in the command file:

```
grant access=BASE to accgroup UnknwGrp for ALL devs
```

To commit the file (`grantbaserights.cmd`) BASE permissions to a non-registered user, enter:

```
symacl -sid 12345 commit -file grantbaserights.cmd
```

Remove permissions from access group

Syntax

In the command file, use the following syntax:

```
remove access=AccessType from accgroup GroupName for accpool PoolName | ALL|NON-POOLED devs;
```


To commit removing permissions, use the following syntax:

```
symacl -sid SymID commit -file FileName
```

Obtain access control information

Only ADMIN or ADMINRD permissions allow viewing the access objects (groups, pools, ACLs). Using the list action without administrative positions, display only the access objects associated with the access group for the host executing the `symacl list` command.

List access control information

Syntax

To list information about groups, pools, and ACLs, use the following syntax:

```
symacl [-sid SymmID|ALL] [-h]  
list [-v]  
list [-accpool | -accgroup | -acl]
```

Examples

To list the access groups on array 0133, enter:

```
symacl -sid 0133 list -accgroup
```

Show access control information

Syntax

To show detailed information about a specified group or pool, use the following syntax:

```
symacl [-sid SymmID | ALL]  
show accpool PoolName [-acl]  
show accgroup GroupName [-acl]
```

 **NOTE:** When using the show actions without administrative permissions, you only see access objects that are associated with the access group to which your host belongs.

Examples

To show the details for access group `ProdB` on array 0133:

```
symacl -sid 0133 show accgroup ProdB
```

Obtain host access ID

Options

`-unique`

Returns the access ID in a segmented, 24-digit numeric form (`xxxxxxxx-yyyyyyyy-zzzzzzzz`).

For example: 12301558-94200021-00347892

Examples

To return the access ID for the controlling host, enter:

```
symacl -unique
```

Verify a locked access control session

Examples

To verify an access control session is locked on any array, enter:

```
symacl list -v
```

Reports the session owner and length of time the session has been locked.

Release locked access control sessions

Description

During the processing of the access control command file, the `prepare` and `commit` actions are critical SYMCLI or SYMAPI operations that are considered access control sessions. In the event a host machine or application should abnormally fail and stop processing any `prepare` or `commit` access operation, the locked session can be aborted.

Syntax

To release the session lock, use the following syntax:

```
symacl release -sid SymmID
```

NOTE: If as a security administrator, you intend to release a lock on a command file session, you must either set the environment variable `SYMCLI_ACCESS_PIN` to your access ID, or enter your PIN every time `symacl` prompts for this.

Access control strategies

This section describes the access control strategies that can be applied in an access-controlled array environment. Several strategies can be considered for establishing or restricting access for a node or group of users or hosts to your environment. These strategies should be considered when setting up an access control environment for the first time.

NOTE: The `discover` command (`symcfg discover`) must be run after making access control changes.

Default configuration: all permissions to users and hosts

The initial (delivered) strategy is to employ a default ID that controls all nodes not yet registered. This default ID can be used to grant a certain level or a minimal level of access for all unregistered nodes.

When an array is delivered it is configured with a group named `UnknwGrp` created for nonregistered hosts (with no ID), as shown in the figure below.



Figure 2. Example default configuration

A special default access ID named *unknown* is added to the group granting all unknown hosts and users ALL permissions. Next, an ACE is created for group *UnknwGrp* granting them ALL permissions to all the array devices. In this scenario, all users and hosts can perform any of the SYMCLI command set operations.

Manage default access

During Access Control implementation, the access IDs of all hosts that need to communicate to the array (for performing array operations using applications such as Solutions Enabler) are registered in the database and granted sufficient permissions to accomplish the functions they need to perform. Once the setup of Access Control is complete, one strategy could be to block or severely limit access for any other host not specifically defined. This is accomplished by adjusting the permissions of the default accgroup *UnknwGrp*.

The default access id *UNKNOWN*, affiliated with the accgroup *UnknwGrp*, is provided at array setup, to allow any host not specifically defined, to have a specific access. The accgroup *UnknwGrp* access rights can be tailored to provide certain hosts a specific access.

The default access id *UNKNOWN* and the accgroup *UnknwGrp* can be removed which will prevent any access by a host, unless it is specifically defined with its own entries. This means that a host, unless specifically defined and associated with a group, will be denied access.

Remove UnknwGrp

Examples

To remove the *UnknwGrp*, enter from the administrative host, enter:

```
delete accgroup UnknwGrp remove_aces=true;
```

Revert UnknwGrp to default

Revert *UnknwGrp* to default

Examples

To revert access control settings back to the default *UnknwGrp*, enter from the administrative host:

```
create accgroup UnknwGrp;  
grant access=ALL to accgroup UnknwGrp for NON-POOLED devs;  
grant access=BASE to accgroup UnknwGrp for ALL devs;  
add default accid name UNKNOWN to accgroup UnknwGrp;
```

Limit access to UnknwGrp group

Examples

To limit access to the *UnknwGrp* with limited access (for example BASE permission only to all undefined hosts), enter:

```
create accgroup UnknwGrp;  
grant access=BASE to accgroup UnknwGrp for ALL devs;  
add default accid name UNKNOWN to accgroup UnknwGrp;
```

Create alternate default ACCID

Description

After removing the default accid *UNKNOWN*, an alternative default accid can be created. A new accgroup is created with permissions that will be the access limitations for all undefined hosts.

Syntax

The default accid provided during setup is called *UNKNOWN*. To create an alternate default accid, use the following syntax:

```
add default accid name IdName to accgroup GroupName
```

Examples

To create an alternate default accid `Others`, enter:

```
create accgroup OtherGrp;  
grant access=BASE to accgroup OtherGrp for ALL devs;  
add default accid name Others to accgroup OtherGrp;
```

These hosts can discover devices, obtain states and statistics from the array, with BASE permission granted to a default accid `Others`.

Establish an administrator

A newly delivered array allows at least one host assigned with administrative (ADMIN) privileges to the access control database.

In this example, to establish an administrator, an access group named `AdminGrp` is created. Then a UNIX workstation named `SunHost1` with an encrypted access ID (of the form `xxxxxxx-yyyyyyy-zzzzzzz`) is added to the `AdminGrp` group. (Obtain an access ID by running the `symacl -unique` command.) Then two ACEs are created: one granting ADMIN permissions and one granting ALL permissions to group `AdminGrp` for all devices in the array.

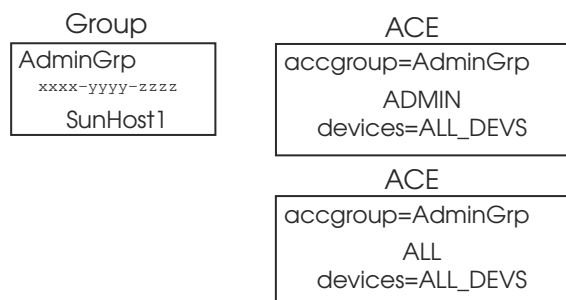


Figure 3. Establish Administrator

Grant all permissions to the nonpooled devices

Once access controlled access pools have been established, access to all the array devices that are not otherwise registered (*Nonpooled*) in any access pool can be set.

In this example, the ACE for the access group named `UnknwGrp` is modified to restrict access to only those devices not registered in an access control pool.



Figure 4. Grant all permissions to non-pooled devices

General absolute access control

General absolute access control registers access only to certain devices on an as-needed basis. In this configuration, the `UnknwGrp` group is removed. Therefore, a node must be known to the array, and only specific users/hosts in defined groups with limited or unlimited permissions have access to certain devices defined in their working pools.

Initial setup summary

Once user needs and limitations are determined, it is highly recommended to use the preview and prepare actions on the first major command file before you committing it; particularly if it is an extensive list. The preview and prepare will identify any coding errors or mistakes in the logic.

Setting up access control for TimeFinder devices

About this task

To set up an access controlled environment for TimeFinder operations, set up both the standard and BCV devices as follows:

Steps

1. Define your working access pool to contain both the standard and BCV SymDevNames.
2. For the group name, grant BASE permissions to access all devices.
3. For the group name, grant BCV permissions to the access pool holding the pairs.

Setting access control for SRDF devices

About this task

To set up an access controlled environment for SRDF operations, set up both the local array and remote array. Because both arrays have their own access controlled database, this requires the following configuration:

Steps

1. With the ADMIN host, create an access control group for the local array. Then with an ADMIN local or remote host, define the same access control group name for the remote array.
2. With an ADMIN host, create an access pool defined with the R1 SymDevnames. Then with the ADMIN local or remote host, define an access pool with the R2 SymDevnames.
3. Grant BASE permissions for the group to access all devices on the R1 array. Then with an ADMIN host, grant BASE permissions for the group to access all devices on R2.
4. Grant SRDF permissions for the group to access the R1 access pool. Then with an ADMIN host, grant SRDF permissions for the group to access the R2 access pool.

Unisphere for VMAX setup strategy

Unisphere software provides a GUI interface to query and manage an array. The GUI is allowed to perform any operations to which the host has been granted rights.

Setup access control for Symaccess

To enable symaccess control for a host, use the permission VLOGIX. This permission must be granted to ALL devices in the array. Note that the behavior of symaccess changes once access pools are created. Once access pools containing devices are created, VLOGIX type privileges are denied to groups unless the privilege is explicitly granted to any group needing access. This includes the default *UnknwGrp*.

Backup and restore the access control database

It is a good practice to create a backup file of the current access control database prior to making changes.

Create an access control database backup file

Syntax

To create a backup of the current access control database, using the following syntax:

```
symacl -sid SymmID backup -file CommandFile
```

The backup operation saves the contents of the access control database in the file specified by the file option. The file must not previously exist. The backup file is compatible for use with the `symacl` command.

NOTE: The backup file contains encrypted versions of the unique IDs. Therefore if comparing the values in the backup file to the original file used to create the database, they will be different.

Restore access control database with backup file

Syntax

To restore the previous configuration data to access control database, use the following syntax:

```
symacl -sid SymmID commit [-v|-noecho] -restore -file CommandFile
```

The restore operation replaces the contents of the access control database with the contents of the file specified by the file option.

NOTE: The backup file contains encrypted versions of the unique IDs, therefore if comparing the values in the backup file to the original file used to create the database, they will be different.

Thin Device Management

This chapter describes thin device management and reporting.

Topics:

- [Thin device management and reporting overview](#)
- [Allocate space on thin devices](#)
- [Reclaim allocations on thin devices](#)
- [Free allocations or free allocations with written tracks](#)
- [Background operations on thin devices](#)
- [Rename thin pools](#)
- [Verify pool and device states](#)
- [Monitor thin pools](#)
- [View all thin device pools](#)
- [View thin device pool details](#)
- [View thin devices](#)
- [View thin device details](#)
- [View thin device allocations bound to a different pool](#)

Thin device management and reporting overview

VMAX3 and VMAX All Flash arrays 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.

Refer to *EMC VMAX3 Family Product Guide for VMAX 100K, VMAX 200K, VMAX 400K with HYPERMAX OS* and *EMC VMAX All Flash Product Guide for VMAX 250F, 450F, 850F, 950F with HYPERMAX OS* for more information on Thin Provisioning.

Use the `symdev`, `symsg` (storage groups), `symcgg` (composite groups), and `symdgg` (device groups) commands to perform the following supported operations for thin devices:

- Create thin devices (refer to [Create devices \(HYPERMAX OS 5977 Q12016SR or higher\)](#) on page 178 or [Create devices \(HYPERMAX OS 5977 lower than 5977 Q12016SR\)](#) on page 180 for creating devices).
- Allocate space on thin devices
- Reclaim space on thin devices
- Free device allocations
- Stop background operations on thin devices
- Rename a thin pool
- Verify pool and device states
- Monitor a thin pool
- View thin devices
- View a thin pool

Allocate space on thin devices

Description

The `allocate` action allocates storage for specified devices. The `-persistent` option specifies that allocated storage cannot be reclaimed or freed. If any of the tracks specified for persistent allocation are already allocated, the already allocated tracks are marked as persistent.

Examples

To start allocation on the thin devices in the device group **myDg**, enter:

```
symdmg -sid 234 -nop -v allocate -dg myDg -persistent
```

The `symmsg` and `symcsg` commands are the same as `symdmg` command using `-sg SgName` and the `-cg CgName`.

To start allocation on thin devices 1078 - 1081, enter:

```
symdev -sid 234 -nop -v allocate -devs 1078:1081 -persistent
```

Restrictions for allocating space on thin devices

The following restrictions apply to the `allocate` action:

- If any of the devices in the device group or storage group are in a state other than allocating or bound, the command will fail.
- If a device group is specified, the action of the command is limited to the standard devices in the device group only.
- Unwritten storage can be recovered with either a `free` or a `reclaim` action. Allocations that were written as zeros can only be recovered by using the `reclaim` action. Written tracks are freed with the `free` action. Additionally, storage that was allocated with the `persistent` attribute can only be recovered by using the `-persistent` option with the `reclaim` action.

Reclaim allocations on thin devices

Description

Reclaiming persistent allocations frees up tracks that are unwritten or zero-based, even if they are marked as persistent. The `reclaim -persistent` command is the only action that frees up persistent tracks and returns a success status if there are no allocations on the specified thin device. If `-persistent` is not specified, this command frees up both unwritten tracks and tracks written with zeros. It will not free up tracks that are marked persistent.

Examples

To start reclaiming space on thin devices in the device group **myDg**, enter:

```
symdmg -sid 234 -nop -v reclaim -dg myDg -persistent
```

The `symmsg` and `symcsg` commands are the same as `symdmg` command using `-sg SgName` and the `-cg CgName`.

To start reclaiming space on thin devices 1078 - 1081, enter:

```
symdev -sid 234 -nop -v reclaim -devs 1078:1081 -persistent
```

Restrictions for using reclaim

- If a device group is specified, the action of the command is limited to the standard devices in the device group only.
- The `reclaim` operation is ignored for any non-thin devices in a range, device group, or storage group.
- If the devices in a device group or storage group are bound to different pools, the `reclaim` operation is allowed on the device group or storage group.
- Space reclamation cannot be performed while a DATA device in the pool is draining.

Free allocations or free allocations with written tracks

Description

To free allocations on thin devices without written tracks, use the `free` command.

To free allocations on thin devices with written tracks, use the `free -all` command. The device must be not mapped or *Not Ready*.

 **CAUTION: Use of `free -all` command can result in lost data. Please use this command carefully.**

Examples

To free allocations on thin devices in device group `myDg`, enter:

```
symdbg -sid 234 -nop -v free -g myDg
```

To free all allocations on thin devices in device group `myDg`, enter:

```
symdbg -sid 234 -nop -v free -g myDg -all
```

The `symmsg` and `symcg` commands are the same as `symdbg` command using `-sg SgName` and the `-cg CgName`.

To free up all allocations on thin devices 1078 - 1081, enter:

```
symdev -sid 234 -nop -v free -devs 1078:1081 -all
```

Restrictions for freeing all allocations

The `free -all` command is blocked on any device participating in replication operations.

If the `free -all` operation is running in the background on a device and it is in a *Ready* state, the following operations are not allowed on that device:

- RDF create pair
- map dev
- convert dev
- All TimeFinder commands

Background operations on thin devices

Description

Actions to `allocate`, `free`, or `reclaim` space on thin devices are performed asynchronously in the background and are started with an explicit start action. Solutions Enabler displays these background tasks as the thin device status, such as *allocating*, *reclaiming*, and so on.

Stop background operations on thin devices

Description

To stop background operations, use the `-stop` option.

Examples

To stop allocation on thin devices in the device group **myDg**, enter:

```
symdmg -sid 234 -nop -v allocate -g myDg -stop
```

The `symmsg` and `symcg` commands are the same as `symdmg` command using `-sg SgName` and the `-cg CgName`.

To stop allocation on thin device 1078, enter:

```
symdev -sid 234 -nop -v allocate -devs 1078 -stop
```

Restrictions

- If a device group is specified, the action of the command is limited to the standard devices in the device group.
- Devices must be in the *allocating*, *reclaiming*, or *freeing* state, otherwise the command fails.

Rename thin pools

Syntax

To rename a thin pool, use the following syntax:

```
rename pool PoolName to NewPoolName
```

Verify pool and device states

Description

The `symcfg verify` command verifies the states of DATA devices and thin devices, and also determines if the pool is in a valid pool state.

To verify if a pool is enabled or disabled, the standard verification options are provided, such as blocking until the pool is in the desired state, and polling at a given rate.

Syntax

To verify a pool state, use the following syntax:

```
symcfg -sid SymmID [-i Interval] [-c Count]
<-pool PoolName |-devs <SymDevStart:SymDevEnd |
SymDevName [,<SymDevStart:SymDevEnd | SymDevName>...]>>
verify -datadev
<-draining | -drainwait | -disabled | -enabled |
-deactivated | -nonpooled | -balancing>
symcfg -sid SymmID [-i Interval] [-c Count]
-pool PoolNameverify -poolstate
<-enabled | -disabled | -balancing>
```

Monitor thin pools

The `symcfg monitor` monitors the thin pools and displays the `-percent` usage.

Syntax

To monitor usage of a thin pool, use the `symcfg monitor` command:

```
symcfg [-sid SymmID] [-i Interval] [-c Count] [-offline]
      [-percent <1-100> -action ScriptFile [-norepeat]]
      [-snap | -rdfa_dse | -thin] [-pool PoolName]
      [-mb] [-gb]
      monitor
```

Options

-action

Selects a script that should be run when the specified percent value is encountered. The full pathname to the action script must be specified. The first argument passed to the script is automatically set to the percent value.

-c

Indicates the number (count) of times to display or to acquire an exclusive lock on the array host database, the local array, and the remote arrays.

-i

Specifies the repeat interval in seconds to display or to acquire an exclusive lock on the array host database, the local , and the remote arrays. The default interval is 10 seconds. The minimum interval is 5 seconds.

monitor

Checks the total percent full of the devices currently configured in the pool and can optionally execute a script file if a specified percentage is encountered.

-norepeat

Specifies that the action script should only be run once if the threshold has been met. Used with the action script option on the `monitor` command.

-percent

Causes the action script to be executed when the percent full argument is encountered.

-pool

Specifies a thin pool.

-thin

Specifies that the pool type is a thin pool.

View all thin device pools

Examples

To view all thin device pools for array 087, enter:

```
symcfg list -pool -all -sid 087
```

Sample output

```
Symmetrix ID: 000197800087

          S Y M M E T R I X   P O O L S
-----
Pool      Flags  Dev      Usable   Free     Used Full  Comp
Name      PTECSL Config  Tracks  Tracks  Tracks (%) Ratio
-----
DG1_FBA_F  TEFEEI RAID-5 (3+1) 27133440 21067637 6065803 22  1.0:1
DG1_FBA_F_4 TEFEEI RAID-5 (3+1) 14283360 3080160 11203200 78  4.0:1
DG1_FBA_F_8 TEFEEI RAID-5 (3+1) 3570840 2048912 1521928 43  2.0:1
Total
Tracks      44987640 26196709 18790931 42  2.0:1
```

Legend:

(P)ool Type:
 S = Snap, R = Rdfa DSE, T = Thin
 (T)echnology:
 S = SATA, F = Fibre Channel, E = Enterprise Flash Drive, M = Mixed, - = N/A
 Dev (E)mulation:
 F = FBA, A = AS400, 8 = CKD3380, 9 or C = CKD3390, - = N/A
 (C)ompression:
 E = Enabled, D = Disabled, N = Enabling, S = Disabling, - = N/A
 (S)tate:
 E = Enabled, D = Disabled, B = Balancing
 Disk (L)ocation:
 I = Internal, X = External, M = Mixed, - = N/A

View thin device pool details

Examples

To view all thin device pool details for pool DG1_FBA_F_8 for array 087, enter:

```
symcfg -sid 087 show -pool DG1_FBA_F_8 -thin -detail
```

Sample output

```
Symmetrix ID: 000197800087

Symmetrix ID           : 000197800087
Pool Name              : DG1_FBA_F_8
Pool Type              : Thin
Disk Location          : Internal
Technology             : EFD
Dev Emulation          : FBA
Dev Configuration      : RAID-5(3+1)
Pool State             : Enabled
Compression State      : Enabled
Compression Ratio      : 2.0:1
# of Devices in Pool   : 13
# of Enabled Devices in Pool : 13
# of Usable Tracks in Pool : 3570840
# of Used Tracks in Pool : 538200
# of Thin Device Tracks : 538200
# of DSE Tracks        : 0
# of Local Replication Tracks : 0
# of Tracks saved by compression : 0
# of Shared Tracks in Pool : N/A
Pool Utilization (%)   : 15
Max. Subscription Percent : N/A
Rebalance Variance     : N/A
Max devs per rebalance scan : N/A
Pool Reserved Capacity : N/A

Enabled Devices(13):
{
-----
Sym      Usable      Free      Used Full FLG Device
Dev      Tracks      Tracks      Tracks (%) S  State
-----
FFEE0    274680    233364    41316  15 .  Enabled
FFEE1    274680    233280    41400  15 .  Enabled
FFEE2    274680    233364    41316  15 .  Enabled
FFEE3    274680    233314    41316  15 .  Enabled
. . .
FFEEC    274680    232910    41770  15 .  Enabled
-----
Tracks    3570840    3032640    538200  15
```

```

}

41316      233364    15 . Enabled
  FFEE1      274680    41400    233280    15 . Enabled
  FFEE2      274680    41316    233364    15 . Enabled
  FFEE3      274680    41366    233314    15 . Enabled
  . . .

```

No Thin Devices Bound to Device Pool DG1_FBA_F_8

Other Thin Devices with Allocations in this Pool (18):

```

{
-----

```

Sym	Bound Pool Name	Total Tracks	Pool Allocated Tracks (%)	Pool Used Tracks
00143	-	46480	23240 50	23240
0014E	-	46780	25779 55	25779
0014F	-	46780	25779 55	25779
. . .				
00500	-	26005	16970 65	16970
Tracks		416455	538200 77	538200

```

}
-----

```

Legend:

Enabled devices FLG:

(S)hared Tracks : X = Shared Tracks , . = No Shared Tracks

Bound Devices FLG:

S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
 D = Deallocating, R = Reclaiming, C = Compressing,
 N = Uncompressing, F = FreeingAll, . = Unbound

View thin devices

Examples

To list thin devices for array 087, enter:

```
symcfg list -tdev -sid 087
```

Sample output

Symmetrix ID: 000197800087

Enabled Capacity (Tracks) : 162267840

Bound Capacity (Tracks) : 396948045

S Y M M E T R I X T H I N D E V I C E S

```

-----

```

Sym	Bound Pool Name	Flgs EMPT	Total Tracks	Total Allocated Tracks (%)	Comp Ratio
00001	-	F..B	26220	10400 40	2.2:1
00002	-	F..B	26220	1333 5	1.0:1
00003	-	F..B	26220	20000 75	1.8:1
00004	-	F..B	26220	0 0	-
00005	-	F..B	26220	3232 12	4.0:1
. . .					
Total Tracks			396948045	60065803 15	2.0:1

```

-----

```

Legend:

```

Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
(M)ultipool : X = multi-pool allocations, . = single pool allocation
(P)ersistent Allocs : A = All, S = Some, . = None
S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
          D = Deallocating, R = Reclaiming, C = Compressing,
          N = Uncompressing, F = FreeingAll, . = Unbound

```

A dash (-) displays in the Bound Pool Name column for all TDEV devices, but only for the first (summary) line for each TDEV.

View thin device details

Examples

To list thin device details for `-devs 140:143` on array 087, enter:

```
symcfg list -tdev -sid 087 -devs 140:143
```

Sample output

Symmetrix ID: 000197800087

```

Enabled Capacity (Tracks) : 162267840
Bound Capacity (Tracks) : 2845170

```

SYMMETRIX THIN DEVICES							
Sym	Bound Pool Name	Flags ESPT	Total Tracks	Pool Subs (%)	Pool Allocated Tracks (%)	Exclusive Allocated Tracks	Comp Ratio
00140	-	F..B	21000	-	0 0	0	-
00141	-	F..B	21000	-	0 0	8000	1.6:1
	DG1_FBA_F	-.--	-	-	5000 24	-	-
	DG3_FBA_F_8	-.--	-	-	16000 76	-	-
00142	-	F..B	21000	-	0 0	0	-
00143	-	F..B	30000	-	0 0	12333	1.8:1
	DG1_FBA_F	-.--	-	-	5000 17	-	-
	DG1_FBA_F_4	-.--	-	-	6000 20	-	-
	DG3_FBA_F_8	-.--	-	-	4000 13	-	-
Total Tracks			93000	-	36000 20	20333	1.7:1

Legend:

```

Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
(S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
(P)ersistent Allocs : A = All, S = Some, . = None
S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
          D = Deallocating, R = Reclaiming, C = Compressing,
          N = Uncompressing, F = FreeingAll, . = Unbound

```

For arrays running HYPERMAX OS 5977, a dash (-) displays in the Bound Pool Name column for all TDEV devices, but only for the first (summary) line for each TDEV. If the TDEV has allocations in other VP Pools, the pool names and allocation details for those pools are displayed on subsequent lines.

View thin device allocations bound to a different pool

Examples

To show all allocations for pool DG1_FBA_F_8 on array 432, enter:

```
symcfg show -pool DG1_FBA_F_8 -thin -sid 432 -all -detail
```

Sample output

When the pool contains allocations from devices that are bound to a different pool:

```
Symmetrix ID: 000197800087

Symmetrix ID           : 000197800087
Pool Name              : DG1_FBA_F_8
Pool Type              : Thin
Disk Location          : Internal
Technology             : EFD
Dev Emulation          : FBA
Dev Configuration      : RAID-5(3+1)
Pool State             : Enabled
Compression State      : Enabled
Compression Ratio      : 2.0:1
# of Devices in Pool   : 13
# of Enabled Devices in Pool : 13
# of Usable Tracks in Pool : 3570840
# of Used Tracks in Pool : 538200
# of Thin Device Tracks : 538200
# of DSE Tracks        : 0
# of Local Replication Tracks : 0
# of Tracks saved by compression : N/A
# of Shared Tracks in Pool : N/A
Pool Utilization (%)   : 15
Max. Subscription Percent : N/A
Rebalance Variance     : N/A
Max devs per rebalance scan : N/A
Pool Reserved Capacity : N/A

Enabled Devices (13):
{
-----
  Sym      Usable      Free      Used Full FLG Device
  Dev      Tracks      Tracks      Tracks (%) S  State
-----
  FFEE0    274680    233364    41316  15  .  Enabled
  FFEE1    274680    233280    41400  15  .  Enabled
  FFEE2    274680    233364    41316  15  .  Enabled
  FFEE3    274680    233314    41366  15  .  Enabled
  . . .
  FFEEC    274680    232910    41770  15  .  Enabled
-----
  Tracks    3570840    3032640    538200  15
}

No Thin Devices Bound to Device Pool DG1_FBA_F_8

Other Thin Devices with Allocations in this Pool (18):
{
-----
  Sym      Bound      Total      Pool      Pool
  Sym      Pool Name    Tracks      Allocated  Used
           Name                               Tracks (%) Tracks
-----

```

```

-----
00143 -          46480      23240  50      23240
0014E -          46780      25779  55      25779
0014F -          46780      25779  55      25779
. . .
00500 -          26005      16970  65      16970
-----
Tracks          416455      538200  77      538200
}

```

Legend:

Enabled devices FLG:

(S)hared Tracks : X = Shared Tracks , . = No Shared Tracks

Bound Devices FLG:

S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
D = Deallocating, R = Reclaiming, C = Compressing,
N = Uncompressing, F = FreeingAll, . = Unbound

Grouping Devices

This chapter describes the benefits of grouping storage array devices, the types of groups, and how to create and modify them.

Topics:

- [Overview of groups](#)
- [Device groups](#)
- [Storage groups](#)
- [Composite groups](#)
- [GNS repository](#)

Overview of groups

Solutions Enabler uses several types of groups to monitor and control storage arrays. The type of group depends on your storage, host, and application environment. Grouping devices allows for:

- Performing control operations on all devices in a group or on device pairs within a group. Refer to [Device groups](#) on page 81.
- Managing groups of storage volumes that belong to a single array for mapping/masking, virtual LUN technology, FAST, and other base control operations. Refer to [Storage groups](#) on page 95.
- Managing groups of devices spread across multiple local arrays. Refer to [Composite groups](#) on page 121.
- Ensuring remote data consistency. Refer to [SRDF consistency groups](#) on page 123.

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.

Use device groups to identify and work with a subset of available devices; obtain configuration, status, and performance statistics on a collection of related devices; or issue control operations that apply to all devices in the specified device group.

A device group can belong to one or more composite groups. For more information, see [Composite groups](#).

Group name services

In a default array environment, device group and composite group definitions are created through a locally-attached host. Upon creation, the group definition is stored in the host configuration database file. Therefore, only the host that created the group can see the group and control it. To perform control operations from another locally-attached host, the group definition must be manually copied to other hosts.

Group Name Services (GNS) can be enabled to store device and composite group definitions in a shared repository located on each array, which then becomes automatically visible to all locally-attached hosts. This allows all GNS-enabled hosts to see the same group definitions across your environment, while sharing real-time updates to group definitions and configurations made by other hosts. For more information on GNS operations, refer to [GNS repository](#).

Names of device groups and devices

Device groups, as well as the devices in a device group, are assigned names that facilitate reference in a session. Assign a device group name when you create it. The name can have up to 31 characters and must be unique for a given configuration database.

When adding a device to a device group, it is given a logical name. This name allows you to refer to the device independently of its physical device name or array device name. The name can have up to 31 characters and must be unique within the device group. It is known only within the context of the device group to which the device belongs. This logical name must be used with the `LdevName` or `ld` argument in any SYMCLI command.

Device group types

When creating a device group, define it as one of the following types:

- REGULAR
- RDF1 (R1 and concurrent R11 devices)
- RDF2 (R2 and concurrent R22 devices)
- RDF21 (cascaded R21 devices)
- ANY (can contain a device mix of REGULAR, RDF1, RDF21 for an R1 site or RDF2, RDF21, RDF22 for an R2/R22 site)

Device lists

A device group must be defined as type REGULAR, RDF1, RDF2, RDF21, or ANY, and may contain various device lists for standard, BCV, virtual (VDEV), and remote devices. A device is placed into a logical list when added to a device group. The following explains the device list types:

- **Standard device list** — Standard device lists provide a mechanism for grouping standard devices in a device group subject to the following restrictions:
 - SRDF and non-SRDF devices cannot be in the same device group *unless* you specify ANY for the type when creating the device group.
 - All SRDF devices in a given device group must belong to the same SRDF group or if concurrent SRDF, belong to two SRDF groups.
- **TimeFinder BCV device list** — TimeFinder BCV device lists provide a mechanism for associating Business Continuance Volume (BCV) regular devices and RDF1 BCV devices with a device group. BCV control operations can be performed on any BCV pair in the device group. Also, you can perform SRDF operations on just the associated RDF1 BCV devices. Refer to *Dell EMC Solutions Enabler TimeFinder SnapVX CLI User Guide* for more information on adding BCV devices to a device group.
- **TF/Snap virtual device list** — A virtual device is a host-accessible device containing address pointers to the data stored on the source device or a pool of SAVE devices, which indicate where the copy session data is located in the physical storage. Virtual devices (VDEV) paired with standard and BCV devices can be associated with any device group. In addition, TF/Snap control operations can be performed on any virtual device in a device group.
- **TF/Clone target list** — Target lists provide a mechanism for establishing source (SRC) and target (TGT) devices for TF/Clone operations. They can be created for both device groups and composite groups. A target list can contain various types of devices, including STDs, VDEVs, or BCV devices (based on a set of rules discussed in [Device restrictions](#)) and can use those devices as targets in clone operations. Remote target lists can also be created for remote operations.

NOTE:

VDEVs cannot be mixed with other types of devices in any target list. In other words, a target list cannot contain both VDEVs and STDs.

- **Gatekeeper device list** — One or more gatekeeper devices can be associated with a device group. SYMCLI uses the associated gatekeeper to issue requests to the array for control operations on the devices within the specified device group. A standard device can be added to a device group. However, the gatekeeper cannot be added to the device group, only associated with a device group. For more information about associating a gatekeeper device with a group, refer to the *Dell EMC Solutions Enabler Installation and Configuration Guide* .

NOTE:

A BCV device or a RDF2 device cannot be assigned as a gatekeeper, nor can a device that is a member of a device group be defined as a gatekeeper.

Device restrictions

Restrictions on the devices you can add to a device group are based on the device type.

RDF1 and RDF2 type device restrictions

The following restrictions apply when adding an SRDF device to a device group:

- All devices in the device group must be SRDF devices.
- All devices in the device group must be either all source (RDF1 type) or all target (RDF2 type) devices.
- All devices in the device group must have the same SRDF group number.

- When there is a combination of R1 and R11 devices in a device group, the R1 and one of its R11 mirrors must have the same SRDF group number. This also applies for R2 and R22 devices.

RDF21 type device restrictions

The following restrictions apply to all devices added to an RDF21 device group:

- All devices must be R21 STD devices. No mixture of STD device types is allowed.
- All R21 devices in a device group must maintain mirror consistency. All R1 mirrors must have the same SRDF group number and all R2 mirrors must have the same SRDF group number.
- Groups that have the same first SRDF group number must have the same cascaded SRDF group number in a cascaded SRDF configuration.
- The existing rules for adding BCV, VDEV, and TGT devices to RDF21 groups apply.

Virtual device restrictions

To add a device as a VDEV, the device must be defined as a virtual device. Virtual devices are VDEV, RVDEV and Hop-2 VDEV.

Clone target restrictions

STD, BCV and VDEV devices can be added to a target list. The following are the sets of device types allowed in a target list, although devices from only one set of devices is allowed in a given device group's target list at any given time:

- Non-SRDF STDs
- R1 STDs
- R2 STDs
- R1 + Non-SRDF Standards
- R2 + Non-SRDF Standards
- Non-SRDF BCVs
- R1-BCVs
- R2-BCVs
- R1-BCVs + Non-SRDF BCVs
- R2-BCVs + Non-SRDF BCVs
- VDEVs (cannot mix VDEVs with other device types in the list)

By default, the logical device name (LdevName) for devices in a target list is TGTxxx. For devices in the remote target list, the default Ldevname is RTGTxxx.

ANY group type restrictions

The ANY group type allows a mix of non-SRDF and SRDF (R1, R11, R2, R22, and R21) devices in a single device group. It lifts the restrictions pertaining to the type of devices added to the device group but still follows the previous restrictions for SRDF devices. For example, all SRDF mirrors in a SRDF list must be of the same device type.

When there is a combination of R1 and R21 devices in a device group, the R1 mirror of the R21 devices must have the same SRDF group number. This also applies to R2 and R21 devices in a device group.

The following devices are allowed in a device group of ANY:

- All REGULAR devices
- All R1 and R11 devices
- All R2 and R22 devices
- All R21 devices
- A combination of REGULAR, R1, R11 and R21 devices
- A combination of REGULAR, R2, R22 and R21 devices

NOTE:

A device group of any SRDF type can change its type because of an `symrdf` control operation. For example, an RDF1 composite group can change to an RDF2 when the device personalities are swapped. The `symrdf swap` operation does not change the type of an ANY device group.

SYMCLI provides various commands to add standard devices, virtual devices, or TF/Clone target devices to a device group. Once a device group is created, SYMCLI commands can be used to add a single device, multiple or all devices on an array, or a list of devices from a file to that group. For more information on adding a list of devices from a file, refer to [Export and import device lists](#).

Create a device group

Create a device group by defining a named empty group of a specific type and then adding devices to the group. A newly-created device group is defined by the devices added to it. Each type of device list has its own set of restrictions.

Create an empty device group

Syntax

To create a device group, use the following syntax:

```
symdbg -type GroupType create GroupName
```

Options

- type

Specifies the group type. Valid group types are: `REGULAR`, `RDF1`, `RDF2`, `R21` or `ANY`. The default type is `Regular`.

name

Specifies the group name. The name can have up to 31 characters and must be unique for a given configuration database.

Example

For example, to create a device group named `prod` whose members are operating as SRDF source devices, enter:

```
symdbg -type RDF1 create prod
```

If a group type is not specified, the default group created is `REGULAR`.

Add devices to a device group

Description

Use the `symdbg` command to add a single device to a device group. Devices can be added by specifying either the physical device name (`add pd`) or the array device name (`add dev`). A logical device name can also be assigned to a device. A valid logical device name cannot exceed 31 characters and must be unique within the device group. If a logical device name is not specified, one will be supplied automatically by `SYMCLI`.

Syntax

The following is the syntax for adding a device:

```
symdbg -g DgName [-offline] [-i Interval] [-c Count] [-v]  
  add pd PdevName [LdevName]  
  
  add dev SymDevName [LdevName] [-sid SymmID]  
    [-rdf | -hop2 | -vdev | -tgt]  
    [-rdfg GrpNum [-remote_rdfg RemoteGrpNum]
```

Options

SymDevName

Specifies the device name when adding virtual devices or target devices.

Interval

The time to wait between attempts to acquire an exclusive lock on the host database.

Count

The number of attempts.

Examples

To add a single device using the physical device name `/dev/rhdisk32` to a device group named **prod**, enter:

```
symdmg -g prod add pd /dev/rhdisk32
```

To add device `00005` to a device group named `prod` and assign the logical device name `temp1`, enter:

```
symdmg -g prod add dev 00005 temp1
```

Add virtual devices to a device group

Syntax

The following is the syntax for adding a virtual device:

```
symdmg add dev SymDevName [LdevName][-sid SymmID]  
[-rdf | -hop2 | -vdev | -tgt]  
[-rdfg GrpNum [-remote_rdfg RemoteGrpNum]
```

Options

-vdev

Specifies that the added device is a virtual device.

-rdf

When adding virtual devices from a remote array (RVDEV), targets the operation to the specified virtual device over SRDF links on the remote array.

-rdfg

With concurrent SRDF, specifies the SRDF group number (*-rdfg GrpNum*), with the *-rdf* option.

-remote_rdfg, -hop2

Targets the operation to the specified virtual device over SRDF links 2 hops away.

Example

For example, to add virtual device `00005` to a device group named **prod1** and assign the logical device name **vdev1**, enter:

```
symdmg -g prod1 add dev 00005 vdev1 -vdev
```

For example, to add virtual device `00005` to a device group named **prod1** from a remote array 2 hops away, enter:

```
symdmg -g prod1 add dev 00005 vdev1 -vdev -hop2 -remote_rdfg 10
```

Add devices to the target list

Description

For TimeFinder/Clone operations, devices can be added to the target device list (TGT) of a device group or the remote target device list (RTGT). STD, SRDF, BCV, and VDEV devices can be added to the TGT, RTGT, and Hop2 TGT target lists.

For details on the types of devices that can be added to a device group's target list, refer to [Clone target restrictions](#) on page 83.

Options

-vdev

Specifies that the added device is a virtual device.

-tgt

Specifies that added devices are from a local array.

-rdf

When adding virtual devices from a remote array (RVDEV), targets the operation to the specified virtual device over SRDF links on the remote array.

-rdfg

With concurrent SRDF, specifies the SRDF group number (`-rdfg GrpNum`), with the `-rdf` option.

-remote_rdfg, -hop2

Targets the operation to the specified virtual device over SRDF links 2 hops away.

Examples

To add target device 00023 to a device group named `prod1` and assign the logical device name `tgt23`, enter:

```
symdmg -g prod1 add dev 00023 tgt23 -tgt
```

To add target device 00069 belonging to SRDF group 12 to a device group named `mywork` and assign the logical device name `tgt2`, enter:

```
symdmg -g mywork add dev 00069 tgt2 -vdev -rdf -rdfg 12
```

Logical device names for virtual devices

When a virtual device is added to a target list, its logical device name is assigned a VDEV default logical name according to these rules:

- VDEVs added to the target list are named TGTxxx.
- VDEVs added to the remote target list are named RTGTxxx.
- VDEVs added to the Hop-2 target list are named 2TGTxxx.

Add ungrouped devices to a device group

Description

The `symdmg addall` command assigns the following logical names to the devices it adds: DEV001, DEV002,..., DEVnnn. Use the `symdmg rename` command to change these logical names after the devices have been added. Or, prior to calling this command, change the default logical device naming conventions using the `SYMCLI_LDEV_NAMING` environment variable. To not truncate logical names too long to fit in the columns of the `symdmg show` and `symcgl show` output, set the `SYMCLI_FULL_LDEVNAME` environment variable.

Syntax

By default, all standard devices (or local virtual devices) are added to a device group, unless options are used to specify certain types of devices. To add multiple devices to a device group, use the following syntax:

```
symdmg -g DgName [-offline] [-i Interval] [-c Count] [-v]
    [-sid SymmID]
    [-vdev | -tgt -rdf | -hop2]
    [-rdfg GrpNum [-remote_rdfg RemoteGrpNum]]
addall dev
    [-SA # | ALL] [-P #] [-N #]
    [-cap # [-captype mb | cyl]]
    [-sel_rdfg SelRdfgNum]
    [-devs SymDevStart:SymDevEnd | SymDevName
    , SymDevStart:SymDevEnd | SymDevName . . .]
```

Options

-i
The time to wait between attempts to acquire an exclusive lock on the host database.

-c
The number of attempts.

-sid
All ungrouped devices from a specific array ID.

-vdev -rdf -rdfg GrpNum
All virtual devices from a local or remote array.

-tgt -rdf -rdfg RemoteGrpNum
All devices are added to the target list of the device group for TF/Clone operations on a remote array.

-vdev -hop2 -remote_rdfg RemoteGrpNum

All virtual devices from a remote array two hops away.

-tgt -rdf -hop2 -remote_rdfg RemoteGrpNum

All devices are added to the target list of the device group for TimeFinder/Clone operations on an array two hops away.

-SA <#|ALL>

All devices visible to one or all front-end directors.

-P #

All devices visible to one or all front-end port number.

-N #

The number of devices to add to the device group.

-devs

Any combination of device ranges and single devices, such as 225:22a,120,a5,7a0:7af

-CAP #

Any combination of device ranges and single devices that are of a specific capacity.

-sel_rdfg

Only SRDF devices belonging to that group number are added.

i **NOTE: When using concurrent SRDF where there are two arrays on the remote side, you must specify the SRDF group number (-rdfg GrpNum) with -rdf option. When the -hop2 option is specified, you must specify a -remote_rdfg RemoteGrpNum.**

Restrictions

The following are not allowed when using `addall` command:

- Mix devices from different arrays.
- Mix SRDF devices that have different SRDF group numbers.
- Add devices defined as a gatekeeper or BCVs.
- Add devices whose device type does not match the device group type.

Set controls on Celerra devices

Syntax

To set controls on Celerra FBA devices with the `syndg -celerra` command, use the following syntax:

```
syndg -g <DgName> [-noprompt] [-i <Interval>] [-c <Count>]
[-bcv | -vdev | -tgt] [-rp] [-star] [-celerra]
    rw_enable [LdevName [LdevName...]] [-p <#>] [-SA <#|ALL>]
    write_disable [LdevName [LdevName...]] [-p <#>] [-SA <#|ALL>]

syndg -g <DgName> [-noprompt] [-i <Interval>] [-c <Count>]
[-bcv [-hop2] | -rbcv | -brbcv | -rrbcv |
-vdev [-hop2] | -rvdev | -tgt [-hop2] |
-rtgt] [-rp] [-star] [-celerra]
    not_ready [LdevName [LdevName...]]
```

List devices in a device group

Syntax

To list devices in a device group, use the following syntax:

```
syndg -g DgName list ld
```

List all device groups

Description

Use the `symdg list` command to list all device groups defined in the configuration database, including group names, group type, array ID, the number of standard, BCV, virtual devices, and TF/Clone target devices (TGTs). The output also indicates if the device group is valid and whether it is contained by a composite group.

Options

`-novalid`

Eliminates the validation of groups during the execution of the `list` command. The `v` column under the `Flags` field and the `v` description in the legend do not display in the output.

Example

```
symdg list
```

NOTE: Device groups with names longer than 17 characters display with their first 17 characters followed by an asterisk (*). [Display composite group information on page 132](#) explains how to set the environmental variable to display longer group names.

Sample output

```
D E V I C E      G R O U P S
                Flags
  Name   Type   VC   Symmetrix ID  Devs   BCVs   VDEVs   TGTs
dgnocgs RDF1   YN   N/A           0       0       0       0
dgingcg REGULAR YY   000194900341 2154   128    0       0
Legend:
Flags:
  V(alid) DG   : Y = Valid, N = Invalid, - = N/A
  (In) C(g)   : Y = Contained by a CG, N = Not contained by a CG
```

Show device group details

Description

Use the `symdg show` command to display information about a specific device group.

Example

```
symdg show dgingcg
```

Sample output

```
Group Name: dgingcg
  Group Type           : ANY
  Device Group in GNS : No
  Valid                : Yes
  Symmetrix ID         : 000194900341
  Group Creation Time  : Tue Dec 1 8:6:9 2009
  Vendor ID            : EMC Corp
  Application ID       : SYMCLI
  Number of STD Devices in Group : 2154
  Number of Locally-associated BCV's : 128
  Number of Locally-associated VDEV's : 0
  Number of Locally-associated TGT's : 0
  Number of Remotely-associated VDEV's (STD RDF) : 0
```



```

Number of Remotely-associated BCV's (STD RDF) : 0
Number of Remotely-associated TGT's (TGT RDF) : 0
Number of Remotely-associated BCV's (BCV RDF) : 0
Number of Remotely-associ'd RBCV's (RBCV RDF) : 0
Number of Remotely-associ'd BCV's (Hop-2 BCV) : 0
Number of Remotely-associ'd VDEV's (Hop-2 VDEV) : 0
Number of Remotely-associ'd TGT's (Hop-2 TGT) : 0
Number of Composite Groups : 2
Composite Group Names : cg110
: cgregular
....

```

Export and import device lists

The list of devices from an existing group can be saved to a file on the host system, and this file can later be imported to create a device group. Device list files are used to recreate a device group that was deleted, or for importing the group into another system.

Export a device list

Syntax

To remove all devices from a group but retain or export the list of devices in the group to a file on your host system, use the following syntax:

```

symdg
    export <DgName> [-delete] [-file <FileName>]
        [[-rdf [-rdfg <GrpNum>]] | [-sid <SymmID>]]
        [-grpfile <GrpDbFileName>]
    exportall [-delete] [-file <FileName>]
        [[-rdf [-rdfg <GrpNum>]] | [-sid <SymmID>]]
        [-grpfile <GrpDbFileName>]

```

Options

- rdf**
Exports an SRDF group. Uses the remote array ID and device names and changes the SRDF group type from R1 to R2 or R2 to R1.
- rdfg GrpNum**
Specifies an SRDF group number.
- delete**
Exports the device group membership to a file and deletes the existing device group in the same operation. To reinstate the same group again, import this list to the same or different device group. Refer to [Import a device list](#) on page 90.

Example

To create a text file that contains the details of all members of the existing device groups, use the `symdg -exportall` operation. To later recreate the device groups from this file, use the `symdg importall` command.

To export the device group membership from group `prod2` to file `prod2list`, enter:

```
symdg export prod2 -f prod2list
```

To export the device group membership to file `prod2list` from group `prod2` and then delete the group, enter:

```
symdg export prod2 -f prod2list -delete
```

For information on deleting a device group, refer to [Delete a device group](#) on page 94.

NOTE:

The `-rdf` option is not supported when exporting R21 device groups.

Import a device list

Description

Typically imported files were previously exported (refer to [Export a device list](#) on page 89).

The `import` action creates the device group if the group name specified in the command does not already exist, devices can be imported to an existing group name that is partially populated. If importing to an existing group, the devices in the imported file are appended to the existing group membership.

To recreate all device groups, use the `symdg importall` command from data contained in a text file that was previously created using the `symdg exportall` command.

NOTE: Raid group members cannot be directly controlled, so exporting or importing details about specific raid group members is not supported.

Syntax

To add multiple devices to a new or existing device group by importing an existing file, use the following syntax:

```
symdg -sid SymmID [-i Interval] [-c Count] [-v]
import DgName [-f FileName]
importall [-f FileName]
```

Example

For example, to create a device group named `prod2` from the file `prod2list`, enter:

```
symdg import prod2 -f prod2list
```

Rename device groups

Description

Use the `symdg rename` command to rename a device group. The new name can contain up to 31 characters and must be unique for the configuration database on which the device group is defined.

Options

-v
For a device group that is a member of one or more composite groups, use the verbose option to view associated composite groups.

Examples

To rename the device group `prod` to `prod_b`, enter:

```
symdg rename prod prod_b
```

To rename a device group that is a member of one or more composite groups, enter:

```
symdg rename prod prod_b -v
```

Sample output

Using verbose option:

```
DG prod contained by CG cg1 was renamed
DG prod contained by CG cg14 was renamed
```

Rename logical device names

Description

Use the `symdmg` command to change the logical name of a device in a device group. The name can have up to 31 characters and must be unique within its device group.

NOTE: This command fails if attempting to rename a logical device name of a device in a device group containing storage groups.

Example

To rename the logical name of device `DEV003` to `TEMP3` in device group named `prod`, enter:

```
symdmg -g prod rename ld DEV003 TEMP3
```

Move a device between device groups

Description

Use the `symdmg` command to move one device from one device group to another. The source and destination device groups must be compatible types.

NOTE: This command cannot be used to move devices from a device group containing storage groups. Use the `symmsg` move or `symmsg moveall` commands to move devices in a storage group.

Syntax

To move an individual device, use the following syntax:

```
symdmg -g DgName [-h] [-offline]
[-i Interval] [-c Count]
move ld LdevName
DestDgName [-force] [-rename]
```

Options

`-i`

Specifies the predetermined time (interval) to wait between attempts to acquire an exclusive lock on the array host database and, for SRDF control operations, on the local and/or remote arrays.

`-c`

Specifies the number of attempts (count).

`-rename`

If there is a device within the destination device group with the same logical device name as device you wish to move into the destination device group, use the `-rename` option to avoid encountering an error. When this option is used, SYMCLI renames the moved device to the next available logical device name as defined in the `SYMCLI_LDEV_NAMING` environment variable.

Example

To move logical device `DEV003` from device group `prod` to device group `test`, enter:

```
symdmg -g prod move ld DEV003 test
```

Move all devices between device groups

Description

Use the `symdg moveall` command to move all standard devices from one device group to another. The source and destination device groups must have compatible device types.

Syntax

To move all devices, use the following syntax:

```
symdg -g DgName [-i Interval] [-c Count] [-v] [-offline]
moveall DestDgName [-force] [-rename]
[-vdev | -tgt [-hop2] | -rvdev | -rtgt]
```

Options

- i**
Specifies the predetermined time (interval) to wait between attempts to acquire an exclusive lock on the array host database and, for SRDF control operations, on the local and/or remote arrays.
- c**
Specifies the number of attempts (count).
- rename**
If there is a device within the destination device group with the same logical device name as the device being moved, use the `-rename` option to avoid an error. When this option is used, SYMCLI renames the moved device to the next available logical device name as defined in the `SYMCLI_LDEV_NAMING` environment variable.
- vdev**
Moves only the virtual devices to the destination device group.
- hop2**
Indicates that the specified device is two hops away.
- tgt**
Moves only the TGT devices to the destination device group.
- rdev**
Moves only remote virtual devices to the destination device group.
- rtgt**
Moves only RTGT devices to the destination device group.

Example

To move all virtual devices from device group `prod` to group `test`, enter:

```
symdg -g prod moveall -vdev test
```

Copy devices between device groups

Description

Devices from an existing device group can be copied into another existing device group of compatible type. Use the `copy` action to copy one standard device from the specified source device group to the destination device group. The source and destination device groups must have compatible types.

Use the `copyall` action to copy all standard devices from the specified source device group to the destination device group. The source and destination device groups must have compatible types. When performing a `copyall` action, the types or number of devices that are included in the copy can be limited using the various filter options.

Syntax

To copy devices from one device group to another device group, use the following syntax :

```
symdbg -g DgName [-i Interval] [-c Count] [-v][-offline]
  copy ld LdevName DestDgName [-force] [-rename]
symdbg -g DgName [-i Interval] [-c Count] [-v]
  [-offline] [-sid SymmID]
  [-SA # | ALL] [-P #] [-N #]
  [-cap # [-capytype mb | cyl]]
  [-sel_rdfg SelRdfgNum]
  [-devs SymDevStart:SymDevEnd | SymDevName
    [, SymDevStart:SymDevEnd | SymDevName...]]
  [-R1 | -R2 | -R21 | -noRDF]
  copyall DestDgName [-force] [-rename]
  [-vdev | -tgt [-hop2] | -rvdev | -rtgt]
```

Copy devices from device group to storage group

Description

Devices from a device group can be copied (or added) to a storage group. The copied devices remain in the device group (*DgName*) and are added to the destination storage group (*SgName*).

If the storage group does not exist, it is created. If optional device types are not specified, only standard devices are added.

NOTE: The `symdbg dg2sg` command is blocked for device groups containing storage groups.

Syntax

To add devices from a device group to a storage group, use the following syntax:

```
symdbg dg2sg
  DgName
  SgName [-bcv | -vdev | -tgt]
```

Example

To add devices from a device group named `prod` to a storage group named `prod_2`, enter:

```
symdbg dg2sg prod prod_2
```

Remove a device from a device group

Description

Use the `symdbg remove` command to remove a device from a device group. By default, the `remove` argument affects standard devices only. However, the `-vdev` option can be specified to remove a virtual device.

NOTE: If you remove the only member of a device group, the device group is not automatically deleted. Use the `symdbg delete` command to explicitly delete the device group, as described in the [Delete a device group](#) on page 94. If the device is contained in a storage group, the operation fails. Use the `symmsg remove dev` command to remove a device in a storage group.

Example

To remove logical device DEV003 from a device group named `prod`, enter:

```
symdmg -g prod -force remove ld DEV003
```

In this example, the `-force` option removes the device regardless of its BCV state.

Remove all devices from a device group

Description

Use the `symdmg rmall` command to remove all devices from a device group. By default, the `rmall` argument affects standard devices only. However, the `-vdev` option can be specified to remove just the virtual devices.

NOTE: Removing all members of a device group, does not automatically delete the device group. Use the `symdmg delete` command to explicitly delete the device group, as described in [Delete a device group](#) on page 94. If the devices are contained in a storage group, the operation fails. Use the `symmsg rmall` command to remove devices in a storage group.

Syntax

To remove all devices from an existing device group, use the following syntax:

```
symdmg -g DgName [-i Interval] [-c Count] [-v]
    [-offline] [-sid SymmID]
    [-SA # | ALL] [-P #] [-N #]
    [-cap # [-capytype mb | cyl]]
    [-sel_rdfg SelRdfgNum]
    [-devs SymDevStart:SymDevEnd | SymDevName
    [, SymDevStart:SymDevEnd | SymDevName...]]
    [-R1 | -R2 | -R21 | -noRDF]
    rmall [-force]
    [-vdev | -tgt -rdf [-rdfg GrpNum] | -hop2]
```

Options

`-SA <#|ALL>`

Specifies devices mapped to a specific front-end (SCSI or Fibre) director number.

`-P #`

Specifies devices mapped to a specific (SCSI or Fibre) director port number.

`-CAP #`

Specifies devices of a specified capacity.

`-devs`

Specifies any combination of device ranges and single devices to remove.

`-vdev | -tgt -rdf [-rdfg GrpNum]`

Specifies local and remote virtual and target devices.

Delete a device group

Syntax

To delete a device group, use the following syntax :

```
symdmg delete DgName [-force]
```

Examples

To delete a device group named `prod`, enter:

```
symdg delete prod
```

When deleting device groups that are associated with a composite group, enter the verbose (`-v`) option to view the composite groups:

NOTE: Deleting populated device groups or a device group of a composite group requires the use of the `-force` option.

```
symdg delete prod -force -v
DG prod was removed from CG cg1
DG prod was removed from CG cg14
```

Perform control operations on device group devices

Description

For additional control operations allowed on device groups refer to Dell EMC Solutions Enabler CLI Reference Guide. Operations can be directed at all devices in a device group. By default the actions only apply to the standard devices in the group. If a control operation is performed on devices in a specific list, the appropriate qualifier needs to be used (such as, `-bcv`, `-rbcv`, etc.).

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. Use the SYMCLI `symmsg` command to create and manage these storage groups.

Storage group restrictions

The following general restrictions apply to all storage groups containing only devices. This applies to storage groups containing child storage groups only, not cascaded storage groups containing both parent storage groups with child storage groups:

- Storage groups must contain only FBA devices, or only CKD devices. A mix of FBA and CKD devices is not allowed.
- Storage groups with CKD devices do not have a Workload.
- GNS does not support storage groups. Storage groups are saved in a special area on the array.
- the maximum number of storage groups for a single array is 16K storage groups.
- Each storage group can contain a maximum of 4096 devices.
- Storage group names can be up to 64 characters in length. Names must begin with an alphanumeric character and may contain embedded hyphens and underscore characters. Names are not case sensitive. Therefore, two storage groups named `test` and `Test` are not allowed.
- Diskless devices are not permitted in storage groups.
- Logical device names are not supported by storage groups.

Additional usage restrictions for storage groups are described in [Add cascaded storage groups](#) on page 105 and [Restrictions for storage groups with defined Host I/O limits](#) on page 108.

Create storage groups

Syntax

To create an empty storage group, use the following syntax:

```
symmsg -sid <SymmID> [-i <Interval>] [-c <Count>] [-v]
      create <SgName>
          [-bw_max <MBperSec>]
          [-iops_max <IOperSec>]
          [-dynamic <NEVER | ALWAYS | ONFAILURE>]
```

```
[-sl <SLName> [-wl <WorkloadName>]
[-srp <SRPName>] [-nocompression]
```

Options

bw_max

Specifies the front-end bandwidth of the devices in the storage group. The valid range for bandwidth is from 1 MB/Sec to 100,000 MBs/sec.

iops_max

Specifies the I/Os per second of the devices over a set of director ports. The valid range for IOPs is from 100 IO/sec to 100,000 IOs/sec but must be specified in units of 100 IO/Sec.

-dynamic

Specifies the Host IO Limit dynamic distribution setting for the storage group as follows:

- NEVER – The Host IO Limits for a storage group are never dynamically redistributed (static).
- ALWAYS – The Host IO Limits for the storage group are always dynamically redistributed.
- ONFAILURE – The Host IO Limits for the storage group are dynamically redistributed only upon failure of a front-end port.

NOTE: Host I/O Limits for storage groups on page 107 describes how to use the **bw_max**, **iops_max**, and **dynamic** options to set Host I/O limits.

-sl

Specifies the Service Level for the storage group as follows in the order of the highest to lowest performance expectation:

- Diamond – emulates EFD performance. The only Service Level supported on All Flash Arrays (VMAX 450K, 850K, and 950K).
- Platinum – emulates between EFD and 15K drive performance.
- Gold – emulates 15K drive performance
- Silver – emulates 10K drive performance
- Bronze – emulates 7.2K drive performance
- Optimized – Balances performance across the whole SRP, based on I/O load, type of I/Os, data pool utilization, and available capacities in the pools. It places the most active data on higher performing storage and least active data on the most cost-effective storage. Optimized Service Level does not use a workload type. If no Service Level is specified then Optimized Service Level is the default for the storage group.

NOTE: CKD devices support only **Diamond**, **Bronze**, and **Optimized Service Level**.

-srp

Specifies a Storage Resource Pool on a storage group.

-wl

Specifies the workload type as follows:

- OLTP – Online Transaction Processing
- DSS – Decision Support System

NOTE: If Workload is not specified then a value of **none** is assigned.

-nocompression (-noc)

When creating a storage group the compression attribute is enabled by default on FAST managed storage groups if the associated SRP supports compression. The compression attribute is removed using this option. Compression is allowed only on VMAX All Flash Array and only FBA devices.

The Service Level and Storage Resource Pool name parameters behave as follows:

- If the *SL Name* and *SRP Name* are not specified, the storage group is created with no Service Level or SRP and it is not FAST-managed.
- If both *SL Name* and *SRP Name*, are specified, the storage group is FAST-managed.
- If only the *SL Name*, is specified, a default Storage Resource Pool for the emulation type of the devices in the storage group is used and the storage group is FAST-managed.
- If only the *SRP Name* is specified, an Optimized Service Level is used with the storage group and it is FAST-managed.

Restrictions

- Requires Storage Admin permission.
- Requires Base access type.
- The command fails if the specified Service Level, workload, or Storage Resource Pool does not exist.
- The command fails if the specified Service Level cannot be supported based on the Storage Resource Pool that is being used by the storage group.
- When setting a workload on an group, the command fails if a Service Level is not set.
- When setting an Service Level, Workload or SRP, the command fails if the device list contains both FBA and CKD devices.
- Workload is not supported for PowerMaxOS.

Create devices and add to storage groups

NOTE: If a storage group is not specified or if the specified storage group is not FAST-managed, the device is created using the default Storage Resource Pool for the device's emulation type and an Optimized Service Level.

Syntax

To add a device to a storage group, use the following syntax:

```
create dev count=<n>,
  size = <n> [MB | GB | CYL],
  emulation=<EmulationType>,
  config=<DevConfig>
  [, preallocate size = <ALL>
  [, allocate_type = PERSISTENT]]
  [, remote_config=<DevConfig>, ra_group=<n>]
  [, sg=<SgName> [, remote_sg=<SgName>]]
  ...
```

Example

To create a device and add it to a storage group named SG_Finance, enter:

```
symconfigure -sid 230 commit -cmd "create dev count = 2, size = 1000 cyl,
emulation = fba, config = TDEV, SG = SG_Finance;"
```

```
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
Establishing a configuration change session.....Established.
Processing symmetrix 000197100230
{
    create dev count = 2, size = 1000 cyl, emulation = fba,
    config = TDEV, SG = SG_Finance;
}
Performing Access checks.....Allowed.
Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

Restrictions

- Storage Admin permission is required.
- The command fails if the specified storage group does not exist.
- The command fails if attempting to add an FBA device to a FAST-managed storage group that contains CKD devices or to add a CKD device to a group that contains FBA devices. This includes previous device create operations that add devices to the same storage group during one configuration change session.
- The command fails if attempting to add devices to a storage group that contains encapsulated devices. This includes previous device create operations that add devices to the same storage group during one configuration change session.

Add existing devices to storage groups

Syntax

To add a device, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
      [-c <Count>] [-v] [-celerra] [-rp] [-ckd]
      add dev <SymDevName>
```

To add multiple devices, or devices in a range or a file, use the following syntax:

```
symmsg -sg <SymDevName> -sid <SymmID> [-i <Interval>]
      [-c <Count>] [-v] [-celerra] [-rp] [-ckd]
      [-SA <# | ALL>] [-p <#>] [-N <#>]
      [-cap <#> [-captype <mb> | <cyl>]]
      [-devs <<SymDevStart>:<SymDevEnd> | <SymDevName>
      [,<<SymDevStart>:<SymDevEnd> | <SymDevName>>...]> |
      -file <DeviceFileName> [-tgt] ]
      addall [pd | devs]
```

Options

-tgt

Specifies adding only target devices listed in a text file. Device text files are either a one or two column format; source devices listed in the first column and target devices listed in second column.

Examples

To add a single device to storage group `prod` on array 123, enter:

```
symmsg -sid 123 -sg prod add dev 30
```

To add all devices that are primarily visible from the host (mapped) to storage group `prod` on array 123, enter:

```
symmsg -sid 123 -sg prod addall pd
```

To add a range of physical devices to storage group `prod` on array 123, enter:

```
symmsg -sid 123 -sg prod addall pd -devs 30:3F
```

To add a range or list of devices to storage group `prod` on array 123:

```
symmsg -sid 207 -sg sg1 addall -devs 64:105,22a,505,600:605,0700
```

To add all devices listed in text file `storgrp_a.txt` to storage group `prod` on array 123, enter:

```
symmsg -sid 123 -file storgrp_a.txt -sg prod addall
```

NOTE: Any device that belongs to storage group that is part of a masking view cannot be added to another storage group.

Restrictions

- Storage Admin permission is required.
- If the storage group is FAST-managed, the command fails if the device is already in another storage group that is FAST-managed.
- If the storage group is FAST-managed, the command fails if adding encapsulated devices.
- The command fails if adding FBA devices and the SG has CKD devices.
- The command fails if adding CKD devices and the SG has FBA devices.
- Workload cannot be specified for CKD devices. The command fails if adding CKD devices to an SG if that SG has Workload set.

Add child storage groups to existing storage groups

Syntax

To add devices or storage groups, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
      [-c <Count>] [-v] [-celerra] [-rp] [-ckd]
      add sg <SgName1>[,<SgName2>,<SgName3>,...,<SgNameN>]]
```

Restrictions

- Storage Admin permission is required.
- The command fails when adding a child storage group to a FAST-managed storage group.
- The command fails if the list of child SGs contain both FBA and CKD devices.
- The command fails if adding child SGs with FBA devices and any current child SG contains CKD devices.
- The command fails if adding child SGs with CKD devices and any current child SG contains FBA devices.
- The command fails if adding CKD devices to an SG if that SG has the Workload set.

Storage group merge operation

Syntax

To merge storage groups, use the following syntax:

```
symmsg -sid <SymmID> -sg <targetSgName> merge <sourceSgName>
```

Options

targetSgName

Specifies the target SG for the merge operation. The target SG can be either a parent SG or a standalone SG. If the target SG is a parent SG the source SG will be added as a child SG. If the target SG is a standalone SG, the devices from the source SG will be merged to the target SG, the source SG and masking view will be deleted.

sourceSgName

Specifies the source SG.

Restrictions

- The source SG cannot be empty.
- The source SG must be a standalone SG.
- Both source and target SG must be in a single masking view with the same IG and PG.

Storage group split operation

Syntax

To split storage groups, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> split <SgName1> -view_name <MvName>
      [-devs <<SymDevStart>:<SymDevEnd> | <SymDevName>
      [,<<SymDevStart>:<SymDevEnd> | <SymDevName>>...]]
```

Options

SgName

Specifies the source SG for the split operation.

view_name <MvName>

Specifies the name of the masking view created as part the split operation.

SgName1

For a cascaded source SG, the <SgName1> option is used to specify the child SG to be split. For the standalone source SG, the <SgName1> option is used to specify the new SG to be created during the split operation.

devs

Specifies the device list to be split from the source SG to the new target SG

Restrictions

- The source SG must be in a single masking view.

Modify storage group properties

Syntax

To modify the storage group properties, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
      [-c <Count>]
      set <[-bw_max <MBperSec> | NOLIMIT ]
          [-iops_max <IOperSec> | NOLIMIT ]
          [-dynamic <NEVER | ALWAYS | ONFAILURE>]>
          [-sl <SL Name> [-wl <Workload Name>] | -nosl]
          [-srp <SRP Name> | -nosrp]
          [-compression | -nocompression]
```

Options

bw_max

Specifies the front-end bandwidth of the devices in the storage group. The valid range for bandwidth is from 1 MB/Sec to 100,000 MBs/sec.

iops_max

Specifies the I/Os per second of the devices over a set of director ports. The valid range for IOPs is from 100 IOs/sec to 100,000 IOs/sec but must be specified in units of 100 IO/Sec.

dynamic

Specifies the Host IO Limit dynamic distribution setting for the storage group as follows:

- NEVER – The Host IO Limits for a storage group are never dynamically redistributed (static).
- ALWAYS – The Host IO Limits for the storage group are always dynamically redistributed.
- ONFAILURE – The Host IO Limits for the storage group are dynamically redistributed only upon failure of a front-end port.



NOTE: Host I/O Limits for storage groups on page 107 describes how to use the bw_max, iops_max, and dynamic options to set Host I/O limits.

-sl

Specifies the Service Level for the storage group as follows in the order of the highest to lowest performance expectation:

- Diamond – emulates EFD performance. The only Service Level supported on All Flash Arrays (VMAX 450K, 850K, and 950K).
- Platinum – emulates between EFD and 15K drive performance.
- Gold – emulates 15K drive performance
- Silver – emulates 10K drive performance
- Bronze – emulates 7.2K drive performance
- Optimized – Balances performance across the whole SRP, based on I/O load, type of I/Os, data pool utilization, and available capacities in the pools. It places the most active data on higher performing storage and least active data on the most cost-effective storage. Optimized Service Level does not use a workload type. If no Service Level is specified then Optimized Service Level is the default for the storage group.

 **NOTE: CKD devices support only Diamond, Bronze, and Optimized Service Level.**

-nosl

Removes Service Level from a storage group.

-wl

Specifies the workload type as follows:

- OLTP – Online Transaction Processing
- DSS – Decision Support System

 **NOTE: If no workload is specified then default workload *none* is set for the storage group.**

-srp

Specifies a Storage Resource Pool on a storage group.

-nosrp

Specifies a Storage Resource Pool to be removed from a storage group.

-compression (-com)

Sets compression on a storage group. Compression is allowed only on VMAX All Flash Array and only FBA devices.

-nocompression (-nco)

Removes compression on a storage group. Compression is allowed only on VMAX All Flash Array and only FBA devices.

Restrictions

- Requires Storage Admin permission.
- Requires Base access type.
- The command fails if the specified Service Level, workload, or Storage Resource Pool does not exist.
- The command fails if the specified Service Level cannot be supported based on the Storage Resource Pool that is being used by the storage group.
- When setting a workload on a group, the command fails if a Service Level is not set.
- When setting a Service Level, Workload or SRP, the command will fail if the device list contains both FBA and CKD devices
- When setting a SRP, the command fails if the SG contains FBA devices and SRP doesn't contain FBA Pools,
- When setting a SRP, the command fails if the SG contains CKD devices and SRP doesn't contain CKD pools.
- When setting a Service Level, the command fails if the SG attached to an SRP contains CKD devices and Service Level is not compatible with CKD emulation.
- When setting a Service Level, the command fails if the SG attached to an SRP contains FBA devices and Service Level is not compatible with FBA emulation.
- Workload cannot be specified for CKD devices. When setting Workload, the command fails if the SG contains CKD devices.

Add Snapshot policy to a storage group

Description

Snapshot policies can be added or removed to/from an SG.

Syntax

To add or remove a Snapshot policy to/from a storage group, use the following syntax:

```
symmsg -sid <SymmID> -sg <SgName> add|remove -policy <PolicyName> -type snap
```

Options

-policy

Specifies the Snapshot policy name that is to be added or removed.

-type

This specifies the type of policy to be added or removed. The value *<snap>* specifies a Snapshot policy.

Set Service Level for a storage group

Description

When setting a Service Level for a storage group, a workload can also be set for storage groups with FBA devices. If no workload is specified, then the workload *none* is assigned.

If the storage group does not have a SRP set, the system default SRP for the emulation type of the devices in the storage group (FBA or CKD) is used and the group becomes FAST-managed.

Syntax

To set a Service Level for a storage group, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> [-sl <SL Name> [-wl <Workload Name>]]
```

Options

-sl

Specifies the Service Level for the storage group as follows in the order of the highest to lowest performance expectation:

- Diamond – emulates EFD performance. The only Service Level supported on All Flash Arrays (VMAX 450K, 850K, and 950K).
- Platinum – emulates between EFD and 15K drive performance.
- Gold – emulates 15K drive performance
- Silver – emulates 10K drive performance
- Bronze – emulates 7.2K drive performance
- Optimized – Balances performance across the whole SRP, based on I/O load, type of I/Os, data pool utilization, and available capacities in the pools. It places the most active data on higher performing storage and least active data on the most cost-effective storage. Optimized Service Level does not use a workload type. If no Service Level is specified then Optimized Service Level is the default for the storage group.

 **NOTE: CKD devices support only Diamond, Bronze, and Optimized Service Level.**

-wl

Specifies the workload type as follows:

- OLTP – Online Transaction Processing
- DSS – Decision Support System

 **NOTE: If no workload is specified then default workload *none* is set for the storage group.**

Remove Service Level from a storage group

Description

When removing the Service Level from a storage group, any workload that was assigned for the Service Level is removed. If the storage group has a SRP, the storage group will be assigned an *Optimized* Service Level. Otherwise, there is no Service Level or SRP for the storage group and the group will no longer be FAST-managed.

Syntax

To remove the Service Level from a storage group, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> -nosl
```

Set SRP for a storage group

Description

When setting a SRP for a storage group. if the storage group does not have an assigned Service Level, an *Optimized* Service Level is assigned and the group becomes FAST managed.

Syntax

To set the SRP for a storage group, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> -srp <SRP Name>
```

Remove SRP from a storage group

When removing an SRP from a storage group, if the storage group has an assigned Service Level, the system default SRP for the emulation type of the devices in the storage group is used. If there is no assigned Service Level, the storage group is no longer FAST-managed.

Syntax

To remove the SRP from a storage group, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> -nosrp
```

Remove Service Level and SRP from a storage group

Syntax

To remove the Service Level and SRP from a storage group, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> -nosl -nosrp
```

 **NOTE:** When the Service Level and SRP are removed from a storage group, it is no longer FAST-managed.

Change workload on a storage group

Description

To change the workload without changing the assigned Service Level, specify both the `-sl` option with the current Service Level name and the `-wl` option with the new workload name.

Syntax

To change the workload type for a storage group, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> -sl <SL Name>  
-wl <Workload Name>
```

Remove workload on a storage group

Description

To remove the workload on a storage group while retaining the current Service Level, specify the `-sl` option with current Service Level name and omit the `-wl` option. This causes the workload **<none>** to be assigned. The storage group will remain FAST-managed.

Syntax

Use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]  
[-c <Count>]  
set <[-bw_max <MBperSec> | NOLIMIT ]  
[-iops_max <IOperSec> | NOLIMIT ]  
[-dynamic <NEVER | ALWAYS | ONFAILURE>]>  
[-sl <SL Name> [-wl <Workload Name>] ]-noslo]  
[-srp <SRP Name> | -nosrp]
```

Restrictions for storage group modification

- Requires Storage Admin permission.
- Requires Base access type.
- The command fails if the user-supplied storage group, Service Level, workload, or Storage Resource Pool does not exist.
- The command fails if the user-supplied Service Level cannot be supported based on the Storage Resource Pool that is being used by the storage group.
- The command fails if setting a Service Level or a Storage Resource Pool to a parent storage group.
- When setting a Service Level or a Storage Resource Pool, the command fails if any device in the storage group is already in another storage group which is FAST-managed.
- When setting a Service Level or a Storage Resource Pool, the command fails if the storage group contains encapsulated devices.
- When setting a workload on a storage group, the command fails if the storage group does not have a Service Level or a Service Level is not being set. It will also fail if a Service Level set on the storage group is being removed.

Standalone storage groups and cascaded groups

Solutions Enabler provides the capability for storage groups to contain other storage groups (cascaded storage groups). This cascading of storage groups allows for individual FAST policies for the storage groups containing devices and a masking view for the storage group containing other storage groups. The storage group, containing only devices, that is contained within a parent storage group is referred to as the child storage group.

Parent and child storage group restrictions

The following restrictions apply to cascaded storage groups. This applies to storage groups containing other storage groups (a parent storage group and children storage groups):

- Only a single level of cascading is permitted. A parent storage group may not be a child of another storage group.
- Storage groups can only contain devices or other storage groups. No mixing is permitted. This covers attempts to add devices to a parent storage group using `add`, `copy`, `move`, and `dg2sg`. This also covers attempts to add child storage groups to an storage group containing devices.
- A parent can have up to 64 child storage groups.
- Empty storage groups can be added to a parent storage group as long as the parent storage group inherits at least one device when the parent storage group is in a view.
- A parent storage group cannot inherit the same device from more than one child storage group.
- A child storage group may only be contained by a single parent storage group.
- No parent storage group can have a FAST association.
- A storage group already associated with a FAST policy is not allowed to be a parent storage group.
- Masking is not permitted for a child storage group which is contained by a parent storage group already part of a masking view.
- Masking is not permitted for the parent storage group which contains a child storage group that is already part of a masking view.
- A child storage group cannot be deleted until it is removed from its parent storage group.

Child storage group operation restrictions

The following restrictions exist for device operations involving child storage group device operations. This includes the `symsg add` and `addall` commands, as well as the `copy`, `copyall`, `move` and `moveall` commands because they involve adding devices to a storage group:

- A `moveall` operation is not permitted from a storage group that contains a masking view or is associated with a FAST policy.
- A `copy` or `copyall` operation is not permitted from a storage group that is associated with a FAST policy into a storage group that is associated with a FAST policy.
- When in a view, the total number of devices inherited by a parent storage group cannot exceed 4096 devices.
- If adding Celerra devices into an storage group or a child storage group with Celerra devices to a parent storage group within a view, you must use the `-celerra` flag.
- If adding a CKD device or a child storage group with CKD devices to a parent storage group within a view, you must use the `-ckd` flag.
- If adding a RecoverPoint tagged device or a child storage group with RecoverPoint tagged devices to a parent storage group within a view, you must use the `-rp` flag.
- Adding an RecoverPoint tagged device to a storage group that is in a masking view containing FCoE directors is not allowed.
- Adding an AS400 device to a storage group that is in a masking view containing FCoE directors is not allowed.

The following restrictions exist for child storage group remove and removeall operations. This also includes move and moveall operations because they perform device removals:

- If the parent storage group has a masking view, any operations involving device removes from the child storage groups will not be permitted, if it causes the parent storage group to have no more devices. This includes removes and moves.

When they are performed on the parent storage group, the following `symmsg` operations affect all of the devices in all of the child storage groups contained by that parent:

- `ready / not_ready`
- `rw_enable / write_disable`
- `hold / unhold`
- `pin / unpin`

Add cascaded storage groups

Description

Use the `symmsg add sg` command to add child storage groups individually to a parent storage group.

NOTE: The `symmsg add` command allows a storage group with Host I/O limits set to be added to a parent storage group that is in a provisioning view using a port group that has FCoE ports.

Syntax

To add child storage groups to a parent storage group, use the following syntax:

```
symmsg -sg SgName -sid SymmID [-i Interval]
        [-c Count] [-v] [-celerra] [-rp] [-ckd]
        add sg SgName1[, SgName2, SgName3, ..., SgNameN
```

Restrictions

The following additional restrictions apply to `symmsg add sg` operations:

- A device cannot be added to a storage group associated with a FAST policy if the device already exists in another storage group that is also associated with a FAST policy.
- Any attempt made to add a storage group that is part of a masking view to a second storage group fails.
- If the Host I/O limits (either `-bw_max` or `-iops_max`) on the child storage group that is being added is greater than what is configured on the parent storage group, the operation will fail.

Convert standalone storage group to cascaded group

Description

Use the `symmsg convert -cascaded` CLI command to non-disruptively convert from a standalone storage group to a cascaded storage group consisting of a parent storage group and a single child storage group. If the standalone storage group has a Host IO Limit it must be specified, if after the conversion the limit will be set on the parent or the child storage group. The standalone storage group has Host I/O Limits set if a limit to either `bw_max` or `iops_max` was configured on the group.

Syntax

To convert a storage group to a cascaded group, use the following syntax:

```
symmsg -sid <SymmID> [-i <Interval>] [-c <Count>] [-v]
        convert -cascaded <SgName> <ChildSgName>
        [-host_IO <on_parent | on_child>]
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: BASE (if the storage group is not in a masking view)
 - Required Authorization Rights: VLOGIX (if the storage group is in a masking view)
- The user-supplied storage group must exist.

- The supplied child storage group must not already exist.
- The supplied storage group must be standalone.
- If the supplied storage group has a Host IO limit defined, the **host_IO** option must be specified .

Convert cascaded group to standalone storage group

Description

Use the `symmsg convert -standalone` CLI command to non-disruptively convert from a cascaded storage group consisting of a parent storage group and a single child storage group to a standalone storage group. If either the parent or the child storage group has a Host IO limit defined, it will be set on the standalone storage group. But if both parent and child storage groups have a Host IO limit, the **host_IO** option must be supplied. A storage group has Host I/O Limits set if a limit to either `bw_max` or `iops_max` was configured on the group.

Syntax

To convert a cascaded group to a standalone storage group, use the following syntax:

```
symmsg -sid <SymmID> [-i <Interval>] [-c <Count>] [-v]
      convert -standalone <SgName> [-host_IO <keep_parent | keep_child>]
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: BASE (if the SG is not in a Masking View)
 - Required Authorization Rights: VLOGIX (if the SG is in a masking View)
- The command will fail if the user-supplied storage group does not exist.
- The command will fail if the supplied storage group is either standalone or a child storage group is in a cascaded storage group configuration.
- The command will fail if the supplied storage group is a parent storage group that contains more than one child storage group.
- The command will fail if both the supplied storage group and its child storage group have a Host I/O Limit defined and the **host_IO** option was not given.

Remove devices from a storage group

Description

Devices can be removed from a storage group as a single device, using a combination of device ranges and single devices, or grouped in a text file.

Options

-tgt Specifies adding only target devices listed in a text file. Device text files are either a one or two column format; source devices listed in the first column and target devices listed in second column.

Syntax

To remove a single device from a storage group, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
      [-c <Count>] [-v] [-celerra] [-rp] [-ckd]
      remove dev <SymDevName> [-force]
```

To remove multiple devices, devices in a range, or a file, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
      [-c <Count>] [-v] [-celerra] [-rp] [-ckd]
      [-SA <# | ALL>] [-p <#>] [-N <#>]
      [-cap <#> [-captype <mb> | <cyl>]]
      [-devs <<SymDevStart>:<SymDevEnd> | <SymDevName>]
```

```
[,<<SymDevStart>:<SymDevEnd> | <SymDevName>>>...]> |  
-file <DeviceFileName> [-tgt] ]  
rsmall [-force]
```

Example

To remove all devices that are listed in text file `storgrp_a.txt` from storage group `prod` on array 123, enter:

```
symmsg -sid 123 -file storgrp_a.txt -sg prod rsmall
```

NOTE: Any attempt made to remove a device belonging to a storage group that is part of a masking view to a second storage group fails.

Remove child storage groups

Description

To remove child storage groups individually from a parent storage group, use the `symmsg remove sg` command.

Syntax

Use the following syntax to remove child storage groups from a parent storage group:

```
symmsg -sg SgName -sid SymmID [-i Interval]  
      [-c Count] [-v] [-celerra] [-rp] [-ckd]  
      remove sg SgName1 [, SgName2, SgName3, SgNameN]
```

Restrictions

The following restrictions apply to `symmsg remove sg` operations:

- If the parent storage group has a masking view, any operations involving device removes from the child storage groups will not be permitted, if it causes the parent storage group to have no more devices. This includes removes and moves.
- You cannot remove devices from a child storage group using the parent storage groups name.
- If you are removing Celerra devices from a storage group or a child storage group with Celerra devices from a parent storage group within a view, you must use the `-celerra` flag.
- If you are removing a CKD device or a child storage group with CKD devices from a parent storage group within a view, you must use the `-ckd` flag.
- If you are removing an RecoverPoint tagged device or a child storage group with RecoverPoint tagged devices from a parent storage group within a view, you must use the `-rp` flag.

Host I/O Limits for storage groups

The VMAX array Host I/O Limits feature allows you to define limits to enforce service levels and make application performance more predictable. The Host I/O Limits settings (`-bw_max MBperSec` and `-iops_max IOperSec`) allow you to limit front-end (FE) port performance by setting FE bandwidth limits on a storage group. This feature is used to limit the amount of FE bandwidth and I/Os per second (IOPs) that can be consumed by a set of devices over a set of director ports. The bandwidth and I/Os controls are then monitored by the VMAX array to ensure that they do not exceed the specified maximum bandwidth or maximum IOPs. This feature allows you to place limits on the FE bandwidth and IOPs consumed by applications on the array.

Host I/O Limits can be added, removed, or modified for a storage group. The Host I/O Limit for a cascaded storage group can be added for both the parent and the child storage group. If a parent storage group has a control set, the setting is shared among all its child storage groups when a provisioning view is created using the parent storage group. If a parent storage group has a control set, you cannot create provisioning views using the child storage groups.

If the Host I/O Limits (either `-bw_max` or `-iops_max`) being configured on a child storage group is greater than what is set on the parent storage group, then the operation will fail. If a Host I/O Limits (either `-bw_max` or `-iops_max`) being configured on a parent storage group is less than the what is set on any one of its child storage groups, then the operation will fail.

NOTE: For additional documentation on using this feature, refer to the *Host I/O Limits for Symmetrix Family Arrays Technical Notes*, which are available on Dell EMC Online support at: <https://support.EMC.com>. The paper explains the benefits of implementing the Host I/O Limits feature and provides use case examples.

Restrictions for storage groups with defined Host I/O limits

The following restrictions apply to a storage group that has defined Host I/O Limits:

- At any given time, a storage group with defined Host I/O Limits can be associated with at most, one port group in any provisioning view. This means, that if the storage group with defined Host I/O Limits is in a provisioning view with a port group, the storage group and port group combination must be used when creating other provisioning views on this storage group. If you attempt to create the view using a different port group, the following error returns:

```
The operation cannot be performed because the storage group with a Host I/O Limit can be associated with at most one port group in any masking view.
```

- Creating a provisioning view on child storage group is not allowed if the parent storage group has defined Host I/O Limits. If you attempt to create a view, the following error returns:

```
The operation can not be performed because the child storage group or the parent storage group has a Host I/O Limit defined.
```

- Setting Host I/O Limits on a parent storage group is not allowed if any child storage group is already part of a provisioning view that was created using the child storage group. If you attempt to create a view, the following error returns:

```
Cannot perform the requested operation because the group is currently within a masking view or a masking view through cascading.
```

- Any device can be in at most one storage group with Host I/O Limits. If you attempt to add the same device to another storage group with Host I/O Limits, the following error returns:

```
The operation cannot be performed because the device already exists in a storage group with a Host I/O Limit.
```

- If a device is in a masking view with Host I/O Limits, it cannot be in another masking view without Host I/O Limits. If you attempt to add a device in a masking view with Host I/O Limits to another masking view without Host I/O Limits, the following error returns:

```
The operation cannot be performed because the device already exists in a masking view with or without Host I/O Limit.
```

- If the Host I/O Limits (either **-bw_max** or **-iops_max**) being configured on a child storage group is greater than what is set on the parent storage group, then the operation will fail.
- If the Host I/O Limits (either **-bw_max** or **-iops_max**) on the child storage group that is being added is greater than what is configured on the parent storage group, then the operation will fail.
- If a Host I/O Limits (either **-bw_max** or **-iops_max**) being configured on a parent storage group is less than the what is set on any one of its child storage groups, then the operation will fail.
- The Host I/O Limits dynamic (**-dynamic**) setting for the parent and children must match. Child storage groups inherit the dynamic distribution setting from the parent storage group if set.

Set the Host I/O Limit

Description

The **symsg create** command provides two options for setting the storage group Host I/O Limit (**-bw_max MBperSec** and **-iops_max IOperSec**). To configure Host I/O Limits on a set of devices, the front-end limits are added to a storage group. When you create a provisioning view using that storage group, the limits are applied to the devices in the storage group for the ports defined in the port group.

- The **-bw_max** option specifies the front-end maximum bandwidth in MBs/sec for the storage group. The valid range for bandwidth is from 1 MB/Sec to 100,000 MB/Sec.
- The **-iops_max** option specifies the front-end maximum IOs/sec. The valid range for IOPs is from 100 IO/Sec to 2,000,000 IO/Sec and must be specified in units of 100 IO/Sec. An error returns if an invalid range is specified.

Example

To create an empty storage group named **prod** with a front-end bandwidth limit of 40,000 MB per second on array ID 123, enter:

```
symmsg -sid 123 create prod -bw_max 40000
```

Change the Host I/O Limits

Description

The same range limits described in [Set the Host I/O Limit](#) on page 108 apply, but if **NOLIMIT** is specified, then the maximum bandwidth or IOs/Sec is set to unlimited. [List storage groups](#) on page 110 provides more information on listing storage group information including Host I/O Limits feature and demand reports.

NOTE: The **symmsg set** command allows Host I/O Limits to be set on a storage group that is in a provisioning view using a port group with FCoE ports.

Syntax

Use the following syntax to set or change performance limits for an existing storage group:

```
symmsg -sg SgName -sid SymmID [-i Interval] [-c Count]
      set [-bw_max MBperSec | NOLIMIT] [-iops_max IOperSec | NOLIMIT]
```

Example

To set an IOPs limit of 50,000 I/Os per second for a storage group named **prod** on array ID 123, enter:

```
symmsg -sg prod -sid 123 -sg prod set -iops_max 50000
```

Set dynamic distribution for FE Host I/O limits

Description

You can optionally configure a dynamic distribution of the FE Host I/O limits by using the **-dynamic** option with the **symmsg set** and **symmsg create** commands. If the option field is not specified, a default of **NEVER** is used.

Setting port failure capability causes the fraction of the configured Host I/O limits available to a configured port to be adjusted based on the number of ports that are currently online. Setting dynamic distribution causes the configured limits to be dynamically distributed across the configured ports, allowing the limits on each individual port to adjust to fluctuating demand. The dynamic distribution feature is supported on front-end ports from Fibre Channel, iSCSI and FCoE directors.

- If port failure capability is set on the FE Host I/O Limit and one of the directors is taken offline or fails, its percentage of the limit is then redistributed, and the two remaining ports can consume the remaining portion of the limit. Once the faulted director comes back online, the limit is again redistributed across all three ports.
- If dynamic distribution is set on the FE Host I/O Limit, then the maximum IOPs for the provisioned view are distributed dynamically across all three of the active ports within the view. The distribution fluctuates based on the current demand on each of the ports.

NOTE: Use the **symmsg show** command to display the value of a storage group's dynamic distribution setting. If a provisioning view has not been created on the parent storage group, the dynamic distribution displays as **N/A**.

Syntax

Use the following syntax to set or change dynamic distribution for an existing storage group:

```
symmsg -sid <SymmID> [-i <Interval>] [-c <Count>] [-v]
      create <SgName> [-bw_max <MBperSec>]
                   [-iops_max <IOperSec>]
                   [-dynamic <NEVER | ALWAYS | ONFAILURE>]
      . . .
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
      [-c <Count>]

      set [<-bw_max <<MBperSec> | NOLIMIT>]
```

```
[-iops_max <<IOperSec> | NOLIMIT>]  
[-dynamic <NEVER | ALWAYS | ONFAILURE>]>
```

Options

-dynamic

- ALWAYS: indicates that the Host IO Limits for the storage group should always be dynamically redistributed
- NEVER: indicates that the Host IO Limits for the storage group should never be dynamically redistributed
- ONFAILURE: indicates that the Host IO Limits for the storage group should dynamically redistributed only upon a failure of a Front-End Port

When the dynamic distribution attribute is set on a parent storage group, every child storage group inherits the same attribute value. The operation is blocked if you attempt to set the attribute on a child storage group.

NOTE: The parent and child storage groups must have the same dynamic distribution attribute value before they can be placed in a cascaded relationship.

Copy devices from a storage group to a device group

Description

Devices from an existing storage group can be copied (or added) to a device group. The copied devices remain in the storage group (*SgName*) and are added to the destination device group (*DgName*).

Syntax

Use the following syntax to add devices from an existing storage group to a device group:

```
symmsg -sid SymmID [-i Interval] [-c Count] [-v]  
      sg2dg SgName DgName [-bcv | -vdev | -tgt]  
      [-R1 | -R2 | -R21 | -noRDF]
```

This command adds RDF1 (-R1) devices, RDF2 (-R2) devices, RDF21 (-R21) devices, or non-SRDF devices (-noRDF) to the device group. This option must match the device type. If the device group does not exist, an error is returned.

NOTE: For cascaded storage groups, the `sg2dg` command behavior has been modified to create a device group from a parent storage group with the device group containing devices from all the child storage groups.

Example

To copy only R1 devices belonging to group number 008 from storage group `prod` to device group `prod_2`, enter:

```
symmsg sg2dg prod prod_2 -R1 -sel_rdfg 008
```

List storage groups

Description

The `symmsg list` command returns a list of all storage group names and the following details:

- Number of devices
- Number of gatekeepers
- Number of child storage groups — For parent storage groups, displays both the number of child storage groups it contains and the cumulative total of all devices contained in the child storage groups.
- FAST association
- Masking view status
- Cascade status
- Host I/O Limits status
- Compression status

NOTE: Storage groups with names longer than 21 characters display with their first 21 characters followed by an asterisk (*). [Display full group names on page 26](#) explains how to set the environmental variable to display full group names.

Syntax

To retrieve a list of all storage groups for a specified array, use the following syntax:

```
symmsg list -sid SymmID
```

Example

To display all storage groups for array 230, enter:

```
symmsg -sid 230 list
```

NOTE: For the (M)asking View flag, if a parent storage group contains a masking view, all its child storage groups display that they are contained within the same masking view.

Sample output

```
S T O R A G E   G R O U P S
Symmetrix ID:      000197100230
Storage Group Name  Flags      Number   Number Child
                   EFMSLC   Devices  GKs      SGs
-----
AcctPay             FX...X     250      13       0
Customer            AX...X     38        2       0
Ordering            MX...X     19        1       0
Payroll             FX....     28        1       0
TestOrdering        .X....     10        0       0
Testsg              .....     10        0       0
Legend:
Flags:
Device (E)mulation  A = AS400, F = FBA, 8 = CKD3380,
                   9 = CKD3390, M = Mixed, . = N/A
(F)ast              X = Fast Managed, . = N/A
(M)asking View      X = Contained in Mask View(s), . = N/A
Cascade (S)tatus    P = Parent SG, C = Child SG, . = N/A
Host IO (L)imit     D = Defined, S = Shared, B = Both, . = N/A
(C)ompression       X = Compression Enabled, . = N/A
```

Report Host I/O Limit demands

Description

Use the **symmsg list** command to report the Host I/O Limit demand on each individual director port (**-by_port -demand**) or port group (**-by_pg -demand**). The demand reports show that the Host I/O Limit is divided equally among all of the directors in the port group, independent of the number of ports on each director. This means that the demand reports show the same Host I/O Limit on both ports even though the Host I/O Limit is shared by the ports. Because of this, it is recommended that you configure only one of the ports of a director in the same port group.

Only a single front-end emulation of each type (FA, EF, etc.) can be assigned to each director, however each of these emulations can be assigned a variable number of physical ports (up to 32, numbered from 0 - 31). In addition, directors containing a Fibre Channel emulation have 32 virtual ports (numbered 32 - 63), that are reserved for use by internal guests.

Syntax

To list Host I/O Limit demands, use the following syntax:

```
symmsg -sid SymmID [-i Interval] [-c Count] [-V]
```

```
list -by_port -demand [-pg PgName | -dir <# [-p # | ALL] | ALL>]
list -by_pg -demand [-pg PgName
```

Examples

To list a demand report for all director ports, enter:

```
symmsg list -by_port -demand
Symmetrix ID : 000197100001
Director      IO Limit      Bandwidth Limit
-----
DIR:Port      Flags      Maximum Demand  Number  Port Speed  Maximum Demand  Number  Excess
              HD      (IO/Sec)  NOLimit  SGs  (MB/Sec)  (MB/Sec) (%)  NoLimit  SGs  (MB/Sec)
              . . .
Legend:
Flags:
(H)ost I/O Limit Exists      Y = Yes, N = No, M = Mixed, . = N/A
(D)ynamic Distribution      Y = Yes, N = No, . = N/A
. . .
```

- The **Flags H** (Host I/O Limit Exist) column indicates whether a limit demand has been configured on a port, and if storage groups with no limits configured are sharing the port. The **Flags D** (Dynamic Distribution) column indicates if the dynamic option is set.
- The **Maximum Demand** columns indicate the total Host I/O Limit demand on the specified director port in MB/Sec or IO/Sec.
- The **Port Speed** column indicates the bandwidth in MB/Sec for that port (the port negotiated speed).
- The **Number NOLimit SGs** columns indicate the number of storage groups with no limits configured that are sharing the port.
- The **Excess** column indicates the amount of bandwidth in MB/sec that is available on the director port after accounting for the demands on the port.

NOTE: Only storage groups that have a provisioning view created on them are shown as placing a Host I/O Limit demand on the director port. If both the parent and child storage group has a Host I/O defined, the maximum demand is calculated based on the parent Host I/O settings.

To list a demand report for all port groups, enter:

```
symmsg list -by_pg -demand
Symmetrix ID : 000195700123
Port Group    IO Limit      Bandwidth Limit
-----
Name          Host Limit  Maximum Demand  Number  Port Grp  Maximum Demand  Number  Excess
              Exist  (IO/Sec)  NOLimit  Speed  (MB/Sec) (%)  NoLimit  SGs  (MB/Sec)
              . . .
PG_Eng1      Yes          0          1      2000      3000 150      0      -1000
PG_Eng2      Yes         1500         0      2000      3000 150      0      -1000
PG_Eng3      Yes          0          1      1000         0  0          1      +1000
PG_Eng4      No          2000         0      1000      1000 100      0          0
. . .
```

NOTE: Only storage groups that have a provisioning view created on them will be shown as placing a Host I/O Limit demand on the port group.

- The **Host Limit Exist** column indicates whether a limit demand has been placed on a port group, and if storage groups with no limits configured are sharing the port group.
- The **Maximum Demand** columns indicate the total Host I/O Limit demand on the specified port group in MB/Sec or IO/Sec.
- The **Port Grp Speed** column indicates the bandwidth in MB/Sec for that port group (the aggregated port negotiated speed for the ports in the group).
- The **Number NOLimit SGs** columns indicate the number of storage groups with no limits configured that are sharing the port group.
- The **Excess** column indicates the amount of bandwidth in MB/sec that is left available on the port group after the demands have been accounted for.

View storage group details

Examples

To view information about storage group `sg1` on array 087, enter:

```
symmsg show -sid 087 sg1
```

Sample output

This output shows a storage group without a defined Host I/O Limit.

```
Name: sg1

Symmetrix ID           : 000197800087
Last updated at       : Fri Mar 04 10:19:41 2016
Masking Views         : No
FAST Managed          : Yes

Service Level Name    : Diamond
Workload              : <none>
SRP Name              : <none>
VP Saved (%)         : 25.1
Compression Enabled   : Yes (1.8:1)
Compression Ratio     : 1.8:1
Host I/O Limit        : None
Host I/O Limit MB/Sec : N/A
Host I/O Limit IO/Sec : N/A
Dynamic Distribution  : N/A
Number of Storage Groups : 0
Storage Group Names   : N/A
Number of Gatekeepers : 0

Devices (4):
{
-----
Sym      Pdev Name      Device      Attr  Sts      Cap
Dev      Name          Config
-----
000DD   N/A              TDEV                RW     23016
000DE   N/A              TDEV                RW     23016
000DF   N/A              TDEV                RW     23016
000E1   N/A              TDEV                RW     23016
```

Show cascaded storage group details

Examples

To show information for parent storage group `SG_Eng1` for array 601, enter:

```
symmsg show SG_Eng1 -sid 601
```

To show information for child storage group `Eng1_Data` for array 601, enter:

```
symmsg show Eng1_Data -sid 601
```

Sample output

For parent storage group:

NOTE: The Number of Composite Groups and Number of Groups fields display as 0 for a child storage group whose parent belongs to a composite group or device group. Only a storage group that is directly contained by the composite group or device group displays a value.

```
Name: SG_Eng1
Symmetrix ID          : 000195700601
Last updated at      : Wed Apr 11 00:28:16 2017
Masking Views        : Yes
FAST Policy          : No
Host I/O Limit       : Defined
Host I/O Limit MB/Sec : 1000
Host I/O Limit IO/Sec : NoLimit
Dynamic Distribution  : N/A
Number of Storage Groups : 2
Storage Group Names   : Eng1_Data      (IsChild)
                     : Eng2_Data      (IsChild)
Number of Composite Groups : 0
Composite Group Names : N/A
Number of Groups      : 2
Group Names           : DG1
                     : DG2

Devices (20):
{
-----
Sym      Device      Cap
Dev      Pdev Name   Config      Sts      (MB)
-----
1A41    N/A          TDEV        RW       2000
1A42    N/A          TDEV        RW       2000
1A43    N/A          TDEV        RW       2000
1A44    N/A          TDEV        RW       2000
1A45    N/A          TDEV        RW       2000
. . .
```

For child storage group:

NOTE: This example shows a child storage group sharing the Host I/O Limit defined in the parent storage group. This is only shown when a provisioning view is created using the parent storage group. If a provisioning view was not created on the parent storage group, the limit is not reported as being shared.

```
Name: Eng1_Data
Symmetrix ID          : 000195700601
Last updated at      : Wed Apr 11 00:28:16 2017
Masking Views        : Yes
FAST Policy          : No
Host I/O Limit       : Shared
Host I/O Limit MB/Sec : 1000
Host I/O Limit IO/Sec : NoLimit
Dynamic Distribution  : N/A
Number of Storage Groups : 1
Storage Group Names   : SG_Eng1      (IsParent)
Devices (10):
{
-----
Sym      Device      Cap
Dev      Pdev Name   Config      Sts      (MB)
-----
1A41    N/A          TDEV        RW       2000
1A42    N/A          TDEV        RW       2000
1A43    N/A          TDEV        RW       2000
1A44    N/A          TDEV        RW       2000
1A45    N/A          TDEV        RW       2000
. . .
```

For a storage group that is not part of a cascaded relationship, the following fields display:

```
Number of Storage Groups : 0
Storage Group Names      : N/A
```

Export and import device lists

You can export the list of devices from an existing storage group to a text file on your host system, and this file can later be imported to create a storage group. This can be useful if you then delete the group and later wish to recreate it. The device list file contains a device description line for as many devices as there are listed in the storage group. Lines that are blank or include a pound (#) sign in the first column are ignored.

Group files for the **import** and **export** commands contain device parameters in the following format:

```
SymmIDSymDevNameSymDevName. . .
```

Group files for the **importall** and **exportall** commands contain device parameters in the following format:

```
SgNameSymmIDSymDevNameSymDevName. . .
```

Repeat this format in the file for multiple storage groups.

i **NOTE: If a filename is not specified with the `symmsg import` or `export` command, a default file named `symmsg.txt` is created in the directory where the `symmsg` command is executed. If a filename is not specified with the `symmsg importall` or `exportall` command, a default file named `symmsgall.txt` is created in the directory where the command is executed.**

Export a device list

Description

Use the **symmsg export** or **exportall** to export a list of devices from one or all storage groups to a text file to the host system.

The **symmsg export** and **exportall** commands support exporting the Service Level, the workload, and the Storage Resource Pool, set on the storage group. If a storage group has a Service Level or a Storage Resource Pool name set on the storage group, the name is copied to the export file. If an implicitly set Optimized Service Level or the system default Storage Resource Pool for the emulation type is set on the storage group, the string DEFAULT is written in place of the name.

i **NOTE: If the storage group configuration, as specified in the export file, has a Storage Resource Pool set and if the Storage Resource Pool has been renamed between the export and import operations, the storage group will not be imported.**

Syntax

To export a storage group device list, use the following syntax:

```
symmsg -sid SymmID [-i Interval] [-c Count][-v]
  export <SgName> [-file <FileName>] [-offline]
  exportall [-file <FileName>] [-offline]
```

Examples

To export the storage group membership from group `prod2` to device file named `prod2list.txt`, enter:

```
symmsg -sid 123 export prod2 -f prod2list.txt
```

Sample output

The output of the text file with the exported data displays as follows:

i **NOTE: Identifier (s), denotes a child storage group name.**

```
000194900341
S sgchild1
S sgchild2
```

The output of a text file generated from `symmsg exportall` displays as follows:


```
<sgparent>
000194900341
S sgchild1
S sgchild2
<sgchild1>
000194900341
00201
00206
<sgchild2>
000194900341
00404
00405
00406
```

Import a device list

Description

You can add multiple devices to a new or existing storage group by importing an existing file that contains a list of devices created by the **export** action. The **import** action will create the storage group if the group name specified in the command does not already exist, or you can import to an existing group name that is partially populated. If you import to an existing group, the devices in the imported file will be appended to the existing group membership.

In addition, you can recreate all storage groups, using the **symmsg importall** command, that were from data contained in a text file previously created using the **symmsg exportall** command.

 **NOTE:** For **importall**, if any of the storage groups that you are attempting to create already exist, Solutions Enabler displays a message indicating it exists, and the operation will continue to create the next storage group in the list.

The following security privileges are required to execute this command:

- Required Access type: BASE
- Required Authorization Rights: Storage Admin.

The **symmsg import** and **importall** commands support importing devices from the FAST-managed storage group in the import file to a target storage group as follows:

- If the imported storage group exists in the array, the target storage group will have its Service Level, workload, and Storage Resource Pool overwritten with those from the imported storage group. If the imported storage group was FAST-managed, the resulting target storage group will become FAST-managed. If the imported storage group was not FAST-managed, the resulting target storage group will become non FAST-managed.
- If the Service Level set on the storage group does not exist on the target array, the storage group will be imported and set to use an Optimized Service Level with no workload.
- If the workload set on the storage group does not exist on the target array, the storage group will be imported without a workload.
- If the Storage Resource Pool set on the storage group does not exist on the target array, the storage group will be imported and set to the default Storage Resource Pool for the emulation type.
- If the both Service Level and the Storage Resource Pool set on the storage group do not exist on the target array, the storage group will be restored with no Service Level and no Storage Resource Pool and will not be FAST-managed.
- If the resulting target storage group is FAST-managed, the command will fail if any device in the storage group is already in another FAST-managed storage group on the array or if any device in the storage group is an encapsulated device.
- The command will fail if the Service Level used by the storage group cannot be supported based on the Storage Resource Pool that is being used by the storage group.
- The **import** fails if the storage group is configured with a Service Level other than "Optimized" and the SRP set for the storage group is configured for FTS external provisioning. With a Service Level other than "Optimized" and an SRP that is configured for FTS external provisioning **importall** skips the import of these storage groups.

Syntax

Use the following syntax:

```
symmsg -sid SymmID [-i Interval] [-c Count] [-v]
import SgName [-f FileName]
importall [-f FileName]
```

Example

To create a storage group named `prod2` from the device file `prod2list.txt`, enter:

```
symmsg -sid 123 import prod2 -f prod2list.txt
```

Rename a storage group

Syntax

To rename a storage group, use the following syntax:

```
symmsg -sid SymmID [-i Interval] [-c Count]  
      rename OldSgName NewSgName -v
```

NOTE: If the storage group is contained by a composite group or device group, all local device groups containing the storage group update with the new name of the storage group. The update is then propagated to the GNS or RDF daemons and the SYMAPI database.

Example

To rename a storage group named `prod` on array 123 to `prod_B`, enter:

```
symmsg rename -sid 123 prod prod_B
```

Move devices between storage groups

Use the `symmsg move dev` or `symdev moveall` command to move one or all devices from one existing storage group to another existing storage group.

Conditions

Moving a device to another storage group will not disrupt the host visibility for the device if any of the following conditions is met:

- Moving devices between child storage groups of a parent storage group when the view is on the parent storage group.
- Moving devices between storage groups when a view is on each storage group and both the initiator group and the port group are common to the views.
- Moving devices between storage groups when a view is on each storage group and they have a common initiator group. They have different port groups but the same set of ports.
- Moving devices between storage groups when a view is on each storage group and they have a common initiator group. They have different port groups but the target port group is a superset of the source port group (additional ports).

NOTE: If the initiator group has the consistent LUN flag set and Solutions Enabler cannot assign consistent LUNs on the additional ports, the operation will fail.

- The source storage group is not in a masking view.

If none of the conditions are met, the operation is rejected but the move can be forced by specifying the `-force` flag. Note that forcing a move may affect the host visibility of the device.

Move a single device between storage groups

Description

Use the `symmsg move dev` command to move one device, specifying the device name, from one storage group to another storage group. The moved device is deleted from the current storage group (`SgName`) and added to the destination storage group (`DestSgName`). You can specify the interval and count options (`-i` and `-c`) to wait a predetermined time (interval) between attempts (count) to acquire an exclusive lock on the host database.

NOTE: A parent storage group may not be specified as the destination for a `move`, `moveall`, `copy` or `copyall` operation, as a parent storage group may not contain both child storage groups and devices.

The **-force** flag is not required for non-disruptive move operations.

Syntax

Use the following syntax to move an individual device:

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
        [-c <Count>] [-v] [-celerra] [-rp] [-ckd]
. . .
        move dev <SymDevName> <DestSgName> [-force]
```

Example

For example, to move device 30 on array ID#59866000123 from storage group **prod** to storage group **test**, enter:

```
symmsg -sid 123 -sg prod move dev 30 test
```

Restrictions

- The selected or all the devices in a source storage group can be moved to a standalone or child destination storage group that is FAST-managed. The operation will fail if after the move the device is in more than one FAST-managed storage group or if moving encapsulated devices to the storage group.
- If the target SG is FAST Managed, the command fails if moving FBA devices to an SG with CKD devices.
- If the target SG is FAST Managed, the command fails if moving CKD devices to an SG with FBA devices.

Move all devices between storage groups

Description

Use the `symmsg moveall` command to move multiple devices from a storage group in a list, using a combination of device ranges and single devices, or from a text file. The moved devices are deleted from the current storage group (`SgName`) and added to the destination storage group (`DestSgName`).

When moving all devices from one existing storage group to another, you can move all devices, or use the additional options to select specific devices. Specify the interval and count options (**-i** and **-c**) to wait a predetermined time (interval) between attempts (count) to acquire an exclusive lock on the host database.

The `symmsg moveall` command requires the **-force** option under the following conditions:

- Moving devices from a source storage group contained in a masking view, either directly or indirectly (inherited from a parent storage group contained in a masking view) to a destination storage group that is also contained in a masking view, either directly or indirectly.
- Moving devices between two storage groups, both associated with a FAST policy.

Syntax

To move multiple devices, devices in a range, or listed in a file, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
        [-c <Count>] [-v] [-celerra] [-rp] [-ckd]
        [-SA <# | ALL>] [-p <#>] [-N <#>]
        [-cap <#> [-captype <mb> | <cyl>]]
        [-devs <<<SymDevStart:SymDevEnd> | <SymDevName>
        [, <<<SymDevStart>:<SymDevEnd> | <SymDevName>>>...> |
        -file <DeviceFileName> [-tgt] ]
        moveall DestSgName
```

Examples

To move all devices from a storage group named **prod** to a storage group named **test** on array 123, enter:

```
symmsg -sid 123 -sg prod moveall test
```

To move multiple devices from a storage group named **prod** to a storage group named **test** on array 123, enter:

```
symmsg -sid 123 -sg prod moveall -devs 31:35,37,40:43 test
```

Restrictions

- The selected or all the devices in a source storage group can be moved to a standalone or child destination storage group that is FAST-managed. However, the operation fails if after the move the device is in more than one FAST-managed storage group, or if moving encapsulated devices to the storage group.
- If the target SG is FAST Managed, the command fails if moving FBA devices to an SG with CKD devices.
- If the target SG is FAST Managed, the command fails if moving CKD devices to an SG with FBA devices.

Copy a device between storage groups

Move all devices between storage groups

Description

Use the `symmsg copy dev` command to copy one device, specifying the array device name, from one storage group to another storage group. The copied device remains in the current storage group (`SgName`) and is added to the destination storage group (`DestSgName`). Specify the interval and count options (`-i` and `-c`) to wait a predetermined time (interval) between attempts (count) to acquire an exclusive lock on the host database.

Syntax

To copy an individual device, use the following syntax:

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
      [-c <Count>] [-v] [-celerra] [-rp] [-ckd]
      copy dev <SymdevName> <DestSgName>
```

Example

To copy device 30 on array 123 in storage group `prod` to storage group `test`, enter:

```
symmsg -sid 123 -sg prod copy dev 30 test
```

Restrictions

- A device in a source storage group can be copied to a standalone or child destination storage group that is FAST-managed. However, the operation fails if the device is already in another storage group that is FAST-managed or if copying encapsulated devices to the storage group.
- If the target SG is FAST Managed, the command fails if moving FBA devices to an SG with CKD devices.
- If the target SG has is FAST Managed, the command fails if moving CKD devices to an SG with FBA devices.

Copy all devices between storage groups

Description

Use the `symmsg copyall` command to copy multiple devices, from one storage group to another storage group. Multiple devices are copied from a storage group in a list, range, or grouped in a text file. The copied devices remain in the current storage group (`SgName`) and is added to the destination storage group (`DestSgName`).

When choosing to copy all devices from one existing storage group to another, you can copy all devices, or use the additional options to select specific devices. Specify the interval and count options (`-i` and `-c`) to wait a predetermined time (interval) between attempts (count) to acquire an exclusive lock on the host database.

Syntax

To copy multiple devices from a storage group, use the following syntax:

```
symmsg -sg SgName -sid SymmID [-i Interval] [-c Count] [-v]
      [-celerra] [-rp] [-ckd] [-SA <# | ALL>] [-p <#>] [-N <#>]
      [-cap <#> [-captype <mb> | <cyl>]]
```

```
[-devs <SymDevStart:SymDevEnd | SymDevName >
[<,SymDevStart:SymDevEnd | SymDevName...> |
-file DeviceFileName [-tgt]]
copyall DestSgName
```

Examples

To copy all devices from a storage group named `prod` to a storage group named `test` on array 123, enter:

```
symmsg -sid 123 -sg prod copyall test
```

To copy multiple devices from a storage group named `prod` to a storage group named `test` on array 123, enter:

```
symmsg -sid 123 -sg prod copyall -devs 31:35,37,40:43 test
```

To copy multiple devices from device file `test_2.txt` to storage group `test` and destination storage group `test_2` on array 123, enter:

```
symmsg -sid 123 -sg test copyall -file test_2.txt test_2
```

Restrictions

- All the devices in a source storage group can be copied to a standalone or child destination storage group that is FAST-managed. However, the command fails if the device is already in another storage group that is FAST-managed or if copying encapsulated devices.
- If the target SG is FAST Managed, the command fails if moving FBA devices to an SG with CKD devices.
- If the target SG has is FAST Managed, the command fails if moving CKD devices to an SG with FBA devices.

Copy devices from a storage group to a device group

Description

Devices from an existing storage group can be copied (or added) to a device group. The copied devices remain in the storage group (`SgName`) and are added to the destination device group (`DgName`).

If the device group does not exist, it will be created. If optional device types are not specified, only standard devices will be added.

Syntax

Use the following syntax to add devices from an existing storage group to a device group:

```
symmsg -sid
SymmID
[-i Interval] [-c Count] [-v]
sg2dg
SgName DgName [-bcv | -vdev | -tgt]
```

Examples

To add devices from a storage group named `prod` on array 123 to a device group named `prod_2`, enter:

```
symmsg -sid 123 sg2dg prod prod_2
```

To add only the target devices from a storage group named `prod` on array 123 to a device group named `prod_2`, enter:

```
symmsg -sid 123 sg2dg prod prod_2 -tgt
```

 **NOTE:** An error will be returned if the `-tgt` option is specified and the storage group contains both standard and BCV devices.

Delete a storage group

Syntax

Use the following syntax:

```
symsg -sid SymmID [-i Interval] [-c Count] delete SgName [-force] -v
```

NOTE: Use the `-force` option to force the deletion of a storage group that is contained by a device group or composite group. On the host performing the delete, all device groups or composite groups containing the deleted storage group are updated. The update is then propagated to the GNS daemon and the database. Deleting a storage group contained by a composite group causes the RDF daemon to stop monitor the composite group.

Example

To delete an empty storage group named `prod` on array `123`, enter:

```
symsg delete -sid 123 prod
```

Restrictions

- Deleting a storage group that is part of a masking view or associated with a FAST policy is not allowed.
- The following restrictions apply to the `symsg delete` command for cascaded storage groups:
 - A parent storage group cannot be deleted unless the `-force` flag is used.
 - A parent storage group cannot be deleted if it has a masking view
 - A child storage group must be removed from its parent storage group before it can be deleted.

Perform control operations on storage group devices

Description

Use the `symsg -sg` command to perform the following operations on devices within a storage group.

Syntax

Refer to the Dell EMC Solutions Enabler CLI Reference Guide for the `symsg` command syntax and allowable control operations.

Composite groups

A composite group (CG) is a user-defined group comprised of devices or device groups, that can belong to one or more locally-attached arrays and one or more SRDF groups within an array. Members can be individual devices or device groups spanning multiple arrays and SRDF groups.

For example, with a composite group you can make a consistent, restartable local copy of a database that is using volumes spanning two local arrays. Without composite groups, it would be impossible to guarantee that all of the BCVs would be split at the same point in time.

Composite groups are created and managed by using the `symcg` command.

Composite group device members

A single composite group can contain devices from the following different device lists:

- Standard device list (STD) — Non-BCV devices that are local to the host.
- Local Business Continuation Volume (BCV) list — BCV devices local to the host.
- Local VDEV list (VDEV) — Virtual devices that are local to the host.
- Remote VDEV list (RVDEV) — Virtual devices that are remote.
- Remote BCV list (RBCV) — BCV devices that are to be associated with the remote mirrors of the STD devices.

- BCV-Remote BCV list (BRBCV) — BCV devices that are to be associated with the remote mirrors of the local BCV devices.
- Remote-Remote BCV List (RRBCV) — Remote BCV devices that are to be associated with the remote mirrors of the RBCV devices.
- Local TGT list (TGT) — TF/Clone target devices that are local to the host.
- Remote TGT list (RTGT) — TF/Clone target devices that are remote.
- R21 STD devices — R21 devices utilize two mirrors and are considered to be concurrent SRDF devices.

Logical device name support

A composite group can contain logical device names (aliases) for devices. All new composite groups, and any devices added to them, will automatically be assigned an *LdevName* if you do not specify one.

The devices added to device groups and composite groups are assigned a default logical names at add time. The name is unique within the group to which it was added.

Device group members within composite groups

A device group can be a member of more than one composite group.

The integrity of control operations is maintained separately at the device group and composite group level. For example, pairing STD devices with BCV devices is only allowed if both devices are contained within the same device group.

Device group membership of composite groups provides the following:

- Building blocks to creating composite groups.
- Support for auto-correct of the composite group.
- SRDF consistency at the composite group and SRDF group name levels.
- GNS support of the composite group.

Restrictions

Before you create a composite group containing device groups, review the following restrictions:

- GNS does not remotely mirror composite groups containing device groups.
- Composite groups can contain either individual devices or device groups, but not both.
- If any of the device groups contained by a composite group become invalid, the composite group also becomes invalid.
- Cannot add a device group to an enabled composite group.
- A composite group cannot contain multiple device groups that contain the same devices.
- Control operations are only allowed at the device group or composite group level if currently supported at that level. For example, control operations such as `symdg` and `symqos` are allowed at the device group level.
- A device group must exist before it can be added to the composite group.
- The device groups in a composite group must follow the guidelines and restrictions outlined in [Device lists](#) on page 82 and [Create a device group](#) on page 84.

Logical names of devices within a composite group

When adding more than one device group to a composite group, a device group may contain devices with the same logical names, resulting in a composite group having more than one device with the same logical name. To prevent this, the logical name of a device in a device group is not carried over into the composite group. A new logical name is created for each device using the following format:

`xxxxxsssss_ddd`

Where:

- xxxxx is the one of the following reserved words representing the device type:

Table 8. Logical name formats

- | | | |
|---------|---------|---------|
| ○ 2BCV | ○ BRBCV | ○ RTGT |
| ○ 2TGT | ○ DEV | ○ RVDEV |
| ○ 2VDEV | ○ RBCV | ○ TGT |
| ○ BCV | ○ RRBCV | ○ VDEV |

- sssss is the last five characters of the array ID containing the device.

- `dddd` is the device number, in hexadecimal, which is guaranteed to be unique within the array.

These naming conventions guarantee that all devices have unique logical names from the composite group point of view. Devices in the device groups maintain the logical names that were assigned at the time when the devices were added to the composite group.

Composite group types

Use one of the following types to create a composite group:

- REGULAR (Solutions Enabler V7.4 and higher allows the inclusion of SRDF devices)
- RDF1 (R1 and concurrent R11 devices)
- RDF2 (R2 and concurrent R22 devices)
- RDF21 (cascaded R21 devices)
- ANY (can contain as device mix of the above types)

NOTE:

A composite group of any SRDF type can change its type because of a `symrdf` control operation. For example, an RDF1 composite group can change to an RDF2 when the device personalities are swapped. SRDF control operations (such as the `suspend`, `establish`, and `swap`) cannot change the type of an ANY composite group but can affect the devices in that composite group.

SRDF consistency groups

An SRDF consistency group is a composite group comprised of SRDF devices (RDF1, RDF2, or RDF21) acting in unison to preserve dependent write consistency of a database distributed across multiple SRDF systems. If a source R1 device in the consistency group cannot propagate data to its corresponding target R2 device, data propagation from all R1 devices in the consistency group is suspended, halting all data flow to the R2 targets.

Consistency is maintained by using either Multi-Session Consistency (MSC) for SRDF/A or SRDF Engineuity Consistency Assist (RDF-ECA) for SRDF/S. For detailed information about SRDF consistency groups, refer to the Dell EMC Solutions Enabler SRDF Family CLI User Guide.

A group is considered an SRDF consistency group if it is a composite group meeting all of the following criteria:

- Created as type RDF1, RDF2, RDF21, or ANY.
- Contains STD devices.
- Set for consistency using the `-rdf_consistency` option, which registers it with the SRDF daemon, and then enabled using the `symcg enable` command.

NOTE:

Deleting a device group from a composite group enabled for SRDF consistency causes the SRDF daemon to stop monitoring this composite group.

NOTE:

For detailed interoperability information, please refer to E-Lab™ Interoperability Navigator (ELN) which can be reached at <http://elabnavigator.EMC.com>.

Create a composite group

To create a composite group:

1. Define a named empty group of a specific type explained in [Composite group types](#) on page 123.
2. Add devices OR device groups to the composite group.

Create an empty composite group

Description

Use the `symcg` command to create a composite group. When you create a composite group, you assign it a name and a group type.

Syntax

Use the following command syntax to create a composite group:

```
symcgc [-i Interval] [-c Count] [-v]  
create CgName [-type REGULAR | RDF1 | RDF2 | RDF21 | ANY ]  
[-apidb | -rdf_consistency]
```

Options

-i

Specifies the interval to wait between attempts to acquire an exclusive lock on the host database and, for SRDF control operations, on the local and/or remote arrays.

-c

Specifies the number of attempts to acquire an exclusive lock on the host database and, for SRDF control operations, on the local and/or remote arrays.

type

Specifies the group type. If a group type is not specified, the default group created is REGULAR.

Create a composite group from devices in a storage group

Description

Starting with Solutions Enabler V8.0.1, you can add selected members of a storage group to a target composite group. If the composite group does not exist, it is created. If none of the optional device types are specified, the default is to add standard devices.

Syntax

Use the following syntax to create a composite group from specified storage group devices:

```
symmsg -sid <SymmID> [-i <Interval>] [-c <Count>] [-v]  
. . .  
sg2cgc <SgName> <CgName> [-bcv | -vdev | -tgt]  
[-R1 | -R2 | -R21 | -noRDF] [-apidb | -rdf_consistency]
```

Options

sg2cgc

If the storage group is a cascaded storage group, `sg2cgc` creates a composite group for the parent storage group with devices from each of the child storage groups.

Create composite groups with SRDF consistency

Description

Specify the `-rdf_consistency` option parameter with `symcgc create` to register a composite group of type RDF1, RDF2, or RDF21 with the SRDF daemon.

NOTE:

Before the SRDF daemon can begin monitoring and managing a composite group created with SRDF consistency, enable it using the `symcgc enable` command. Note that the `symcgc enable` command will be blocked if the type of the enable being performed is MSC or SRDF-ECA and the scope of the enable contains multiple SRDF groups.

When creating a composite group with SRDF consistency, the composite group name is compared against all existing composite group names for uniqueness. This comparison is not case sensitive, ensuring there are no composite group naming collisions in the Symmetrix File System (SFS).

NOTE:

If devices set for consistency protection are in an existing composite group, you cannot add them to another composite group enabled for consistency protection. However, you can add these devices to a composite group not enabled for consistency protection.

Examples

To create a composite group named **mycg1** of type RDF1 with consistency enabled, enter:

```
symcg create mycg1 -rdf_consistency -type rdf1
```

In the options file, you must set the following option to **ENABLE** to create composite groups with SRDF consistency:

```
SYMAPI_USE_RDFD=ENABLE
```

Set controls on Celerra devices

Syntax

Use the following `-celerra` option to set the `rw_enable`, `write_disable`, `ready`, and `non_ready` controls on Celerra FBA devices in a composite group:

```
symcg -cg CgName [-i Interval] [-c Count]
[-noprompt] [-v] [-force]
[-bcv | -vdev | -tgt] [-star] [-sid SymmID]
[-celerra]
```

Export a composite group to file

Syntax

Use the following syntax to export a composite group to file:

```
symcg [-i Interval] [-c Count] [-v]
export <CgName> [-file <FileName>] [-rdf]
[-grpfile <GrpDbFileName>]
exportall [-file <FileName>] [-rdf]
[-grpfile <GrpDbFileName>]
```

Options

- i**
Specifies an interval to wait between attempts to acquire an exclusive lock on the host database and, for SRDF control operations, on the local and/or remote arrays.
- c**
Specifies the number of attempts to acquire an exclusive lock on the host database and, for SRDF control operations, on the local and/or remote arrays.
- rdf**
If `-rdf` is specified, the remote partners of the STD devices and BCV devices are added to the file instead of those devices that exist in the current local composite group. The RBCV devices will become local BCVs. Non-SRDF BCVs, VDEVs, BRBCVs, and RRBCVs will be ignored in this case. The resulting file will have as many device description lines as the composite group has members.

 **NOTE:**

The `-rdf` option cannot be used on **REGULAR** or **cascaded SRDF** composite groups or when the composite group is of type **ANY**. Specifying the `-rdf` option produces an error if any **Hop-2** devices are detected in the composite group.

Delete a composite group

Syntax

Use the following syntax to delete an existing composite group:

```
symcg [-i Interval] [-c Count] [-v]  
delete CgName [-force] [-symforce]
```

Options

-force

If the composite group has members, the command fails unless `-force` is used. If `-force` is specified, the device members of the group are removed, and the group is deleted.

If the composite group is enabled for SRDF consistency, you must use `-force` to delete it.

Import a composite group

Syntax

Use the following syntax to import a composite group from a previously generated file:

```
symcg [-i Interval] [-c Count] [-v]  
import CgName [-f FileName]  
[-apidb | -rdf_consistency] [-rename]  
importall [-f FileName] [-rdf_consistency]
```

Options

-rdf_consistency

When the `-rdf_consistency` parameter is specified, the composite group is registered with the SRDF daemon.

-rename

If `-rename` is specified, the devices are given the next available default device name when added to the group. If `-rename` is not used, the devices are added with the name specified in the import file. This may result in a failure if a device already has that name in the composite group.

Add standard devices to a composite group

Syntax

Use the following syntax to add a standard device to a composite group:

```
symcg [-i Interval] [-c Count] [-v]  
add pd PdevName [LdevName] -cg CgName
```

or

```
symcg -cg CgName -sid SymmID [-i Interval] [-c Count] [-v]  
[-rdf | -hop2]  
[-rdfg GrpNum [-remote_rdfg RemoteGrpNum]]  
add dev SymDevName [LdevName] [-vdev | -tgt]
```

NOTE:

If the composite group already contains one or more storage groups, the command to add devices will fail as this action violates the restrictions for device groups containing storage groups.

Options

- i**
Specifies an interval to wait between attempts to acquire an exclusive lock on the host database on the local and/or remote arrays.
- c**
Specifies the number of attempts to acquire an exclusive lock on the host database on the local and/or remote arrays.
- sid**
Specifies the array ID when adding devices by specifying a device name.
- rdf**
When specified with the `-vdev` option, adds the devices to the remote virtual device list. When specified with the `-tgt` option, adds the target devices to the remote target device list.
- hop2**
Same as `-rdf`, but when the device is two hops away.
- rdfg**
Specifies the SRDF group number.
Composite groups allow remote BCV devices (RBCVs) on both links to be associated with the SRDF group. These RBCVs may be R1 or R2 type BCVs and may also have a remote BCV (RRBCV) associated behind each.
- vdev**
When the `-vdev` option is specified, the devices are added to the virtual device list.
NOTE: In addition, to add a VDEV to the remote virtual device list, specify the `-rdf` option with the `-vdev` options. If the device is two hops away, specify the `-hop2` option instead.
- tgt**
When the `-tgt` option is specified, the devices are added to the target device list.
NOTE: In addition, to add a target device to the remote target device list, specify the `-rdf` option with the `-tgt` options. If the device is two hops away, specify the `-hop2` option instead.

Move and copy devices in composite groups

Devices from an existing composite group can be moved or copied into another existing composite group of compatible type.

- When the `move ld` or `moveall` action is used, the devices are removed from the source composite group and added to the destination composite group.
NOTE:
The `move` and `moveall` commands fail if you attempt to move devices from a composite group containing storage groups. You must use the `symsg move` or `symsg moveall` commands to move devices in a storage group.
- When the `copy ld` or `copyall` action is used, copies of the devices are added to the destination composite group, and the source composite group remains unchanged.
NOTE:
The `symcg copy` and `symcg copyall` commands do not allow you to copy any devices from a composite group containing device groups. In this case, use the `symdg copy` or `symdg copyall` command to copy devices from a device group.

You can assign a new name to a device being copied or moved.

Syntax

Use the following syntax to move or copy devices from one composite group to another.

```
symcg -cg CgName [-i Interval] [-c Count] [-v]
  move ld LdevName DestCgName [-force] [-rename]
  copy ld LdevName DestCgName [-force] [-rename]
symcg -cg CgName [-i Interval] [-c Count] [-v]
  [-sid SymmID]
  [-SA # | ALL] [-P #] [-N #]
  [-cap # [-captype mb | cyl]]
  [-vdev | -tgt [-hop2] | -rvdev | -rtgt]
  [-rdfg GrpNum [-remote_rdfg RemoteGrpNum]]
  [-sel_rdfg SelRdfGrpNum]
  [-devs SymDevStart:SymDevEnd | SymDevName
    [, SymDevStart:SymDevEnd | SymDevName...]]
moveall DestCgName [-force] [-symforce] [-rename]
  [-R1 | -R2 | -R21 | -noRDF]
copyall DestCgName [-force] [-symforce]
  [-R1 | -R2 | -R21 | -noRDF]
```

Options

-sel_rdfg

When the **-sel_rdfg** option is specified, only SRDF devices belonging to that group number are moved or copied.

-r1

Limits the number of SRDF devices added to a composite group to R1s.

-r2

Limits the number of SRDF devices added to a composite group to R2s.

-r21

Limits the number of SRDF devices added to a composite group to R21s.

-noRDF

By specifying **-noRDF**, you can prohibit any SRDF devices from becoming members of the composite group.

Copy devices from a device group to a composite group

Description

Use the **symcg dg2cg** command to copy or add devices from an existing device group to a composite group. If the composite group does not exist, it is created. By default, all device lists from the device group are added to the composite group.

 **NOTE:** During a device add operation, the command fails if the composite group already contains storage groups.

Syntax

Use the following syntax to add devices from an existing device group to an existing composite group:

```
dg2cg DgName CgName [-rename] [-force]
  [-bcv [-hop2] | -nobcv | -rbcv | -rrbcv | -brbcv |
  -vdev [-hop2] | -rvdev | -tgt [-hop2] | -rtgt]
  [-apidb | -rdf_consistency]
```

Options

-bcv

Adds only the BCVs to the composite group.

-vdev

Adds only the virtual devices to the composite group.

-hop2	Indicates the specified device is two hops away.
-rbcv	Adds only the RBCVs to the composite group.
-rrbcv	Adds only the RRBCVs to the composite group.
-brbcv	Adds only the BRBCVs to the composite group.
-nobcv	Adds only the STDs to the composite group.
-rvdev	Adds only the remote virtual devices to the composite group.
-tgt	Adds only the TGTs to the composite group.
-rtgt	Adds only the RTGTs to the composite group.

Remove a standard device from a composite group

Description

Use the **symcg remove** command to remove devices from an existing composite group. You can either remove devices individually, remove all devices, or remove all devices meeting the specified criteria.

Removing an consistency-enabled device does not disable the device. Use **symcg disable** to disable the composite group. If SRDF consistency is enabled and cannot be disabled, use **-symforce**.

NOTE: If the device is contained in a storage group, the operation will fail. Use the **symsg remove dev** or **symsg rma11** command to remove devices in a storage group.

Syntax

Use the following syntax to remove a standard device from a composite group:

```
symcg -cg CgName -sid SymmID [-i Interval] [-c Count] [-v]
      [-rdf | -hop2]
      [-rdfg GrpNum [-remote_rdfg RemoteGrpNum]]
      remove dev SymDevName [-force] [-symforce]
      [-vdev | -tgt]
```

or

```
symcg -cg CgName [-i Interval] [-c Count] [-v]
      remove ld LdevName [-force] [-symforce]
```

or

```
symcg -cg CgName [-i Interval] [-c Count] [-v]
      remove [pd] PdevName [-force] [-symforce]
```

Remove all standard devices from a composite group

Description

Use the **symcg rma11** command to remove all devices from an existing composite group.

NOTE: If the devices are contained in a storage group, the operation will fail. Use the `symmsg remove dev` or `symmsg rmall` command to remove devices in a storage group.

Syntax

Use the following syntax to remove all standard devices from a composite group:

```
symcg -cg CgName [-i Interval] [-c Count] [-v]
  [-sid SymmID]
  [-SA # | ALL] [-P #] [-N #]
  [-cap # [-captype mb | cyl]]
  [-rdf | -hop2] [-vdev | -tgt]
  [-rdfg GrpNum [-remote_rdfg RemoteGrpNum]]
  [-sel_rdfg SelRdfGrpNum]
  [-devs SymDevStart:SymDevEnd | SymDevName
  [,SymDevStart:SymDevEnd | SymDevName...]]
  rmall [-rdf | -hop2] [-force] [-symforce]
  [-R1 | -R2 | -R21 | -noRDF]
```

Options

`-sel_rdfg SelRdfGrpNum`

Indicates to remove only the devices in the specified group from the composite group

Rename a composite group

Description

Use the `symcg rename` command to rename a composite group.

NOTE: The command fails if you attempt to rename a logical device name of a device in a composite group containing storage groups. Also, you cannot rename an SRDF consistency composite group that is enabled.

Syntax

Use the following syntax:

```
symcg [-i Interval] [-c Count]
  rename OldCGName NewCGName
```

Set SRDF group attributes

Description

Use the `symcg set` command to associate a logical name or an SRDF/Star recovery SRDF group number with an SRDF group to perform operations on all associated devices. If the `-rdfg` argument is not specified, then the action is applied to all SRDF groups.

Syntax

Use the following syntax:

```
symcg -cg CgName [-i Interval] [-c Count] [-v]
  set -name [Name] | -recovery_rdfg GrpNum
  [-rdfg SymmID:GrpNum[,GrpNum,...] | all[,...] |
  name :RdfGroupName[,RdfGroupName]]
```

Composite group creation and output

Description

Use the **symcg list** command to obtain information about the composite groups visible to your host system.

Options

-inactive

In a GNS-enabled environment, returns the list of composite groups from the inactive group definition list. For information on active and inactive group lists in GNS, refer to [Active vs. inactive group lists](#) on page 137

-novalid

Eliminates the validation of groups during the execution of the **list** command. The **valid** column does not display in the output.

Syntax

Use the following syntax to list the host-visible composite groups:

```
symcg list
```

Examples

The following example shows how to create a non-consistent RDF1 composite group, associate devices with the group:

Create a non-consistent RDF1 composite group named **MyNewCG**.

```
symcg create MyNewCG -type rdf1
```

Add RDF1 devices to the composite group.

```
symcg -cg MyNewCG addall -devs CEE:CFD -sid 79  
symcg -cg MyNewCG addall -devs 2:7 -sid 32
```

Associate RBCV devices with the composite group.

```
symbcv -cg MyNewCG associateall -devs E92:EA1 -rdf -rdfg 64 -sid 79  
symbcv -cg MyNewCG associateall -devs 48:4D -rdf -rdfg 1 -sid 32
```

Associate RRBCV devices with the composite group.

```
symbcv -cg MyNewCG associateall -devs 328:337 -rrdf -rdfg 64 -sid 79  
symbcv -cg MyNewCG associateall -devs 30:35 -rrdf -rdfg 1 -sid 32\
```

Sample output

As a result of the above command, the **symcg list** command returns the following output:

```
      C O M P O S I T E      G R O U P S  
  
      Name  Type  Valid  Number of  Number of  
      Name  Type  Valid  Symms RAGs  DGs  Devs  BCVs  VDEVs  TGTs  
MyNewCG RDF21  Yes    1    2    2    3    0    0    0
```

Use the **symcg show** command to return additional details about a composite group.

Display composite group information

Syntax

Use the following syntax to list the available composite groups:

```
symcg [-i Interval] [-c Count] [-v]
  list [-offline] [-v]
symcg list
```

Name	Type	Valid	C O M P O S I T E			G R O U P S			
			Symms	RAGs	DGs	Devs	BCVs	VDEVs	TGTs
cgr21	RDF21	Yes	1	2	2	3	0	0	0
cgreg	REGULAR	Yes	0	0	0	0	0	0	0

i **NOTE:** Composite groups with names longer than 21 characters display with their first 21 characters followed by an asterisk (*). [Display composite group information on page 132](#) explains how to set the environmental variable to display full group names.

Show composite group details

Description

Use the **symcg show** action to return additional details about a composite group.

Example

To show the details about composite group **cgr21**, enter:

```
symcg show cgr21
Composite Group Name: cgr21
Composite Group Type      : RDF21
Valid                     : Yes
CG in PowerPath          : No
CG in GNS                 : No
RDF Consistency Protection Allowed : No
RDF Consistency Mode     : NONE
Concurrent RDF           : YES
Cascaded RDF             : No
Number of RDF (RA) Groups : 2
Number of STD Devices    : 3
Number of CRDF STD Devices : 3
Number of BCV's (Locally-associated) : 0
Number of VDEV's (Locally-associated) : 0
Number of TGT's Locally-associated : 0
Number of CRDF TGT Devices : 0
Number of RVDEV's (Remotely-associated VDEV) : 0
Number of RBCV's (Remotely-associated STD-RDF) : 0
Number of BRBCV's (Remotely-associated BCV-RDF) : 0
Number of RRBCV's (Remotely-associated RBCV) : 0
Number of RTGT's (Remotely-associated) : 0
Number of Hop2 BCV's (Remotely-assoc'ed Hop2 BCV) : 0
Number of Hop2 VDEV's (Remotely-assoc'ed Hop2 VDEV) : 0
Number of Hop2 TGT's (Remotely-assoc'ed Hop2 TGT) : 0
Number of Device Groups : 2
Device Group Names      : dg1
                        : dg2

Number of Symmetrix Units (1):
  1) Symmetrix ID      : 000194900341
     Microcode Version : 5977
     Number of STD Devices : 3
     Number of CRDF STD Devices : 3
     Number of BCV's (Locally-associated) : 0
     Number of VDEV's (Locally-associated) : 0
```

```

Number of TGT's Locally-associated          : 0
Number of CRDF TGT Devices                  : 0
Number of RVDEV's (Remotely-associated VDEV) : 0
Number of RBCV's (Remotely-associated STD RDF) : 0
Number of BRBCV's (Remotely-associated BCV-RDF) : 0
Number of RTGT's (Remotely-associated)       : 0
Number of RRBCV's (Remotely-associated RBCV) : 0
Number of Hop2BCV's (Remotely-associ'ed Hop2BCV) : 0
Number of Hop2VDEVs (Remotely-associ'ed Hop2VDEV) : 0
Number of Hop2TGT's (Remotely-associ'ed Hop2TGT) : 0
Number of RDF (RA) Groups (2):
  1) RDF (RA) Group Number                  : 11          (0A)
     Remote Symmetrix ID                    : 00019490237
     Microcode Version                      : 5977
     Recovery RA Group                      : N/A         (N/A)
     RA Group Name                          : N/A
     STD Devices (3):
-----
LdevName      PdevName      Sym  Device      Flags  Cap
Dev           Config     Dev  Config     Sts CSRT  (MB)
-----
DEV341_30     /dev/sd1      0030 RDF21+R-5  WD  X--2  2063
DEV341_31     /dev/sdm      0031 RDF21+R-5  WD  X--2  2063
DEV341_E6     N/A           00E6 RDF21+R-5  WD  X--2  2063
  2) RDF (RA) Group Number                  : 18          (11)
     Remote Symmetrix ID                    : 00019490016
     Microcode Version                      : 5876
     Recovery RA Group                      : N/A         (N/A)
     RA Group Name                          : N/A
     STD Devices (3):
     CRDF STD Devices (3):
-----
LdevName      PdevName      Sym  Device      Flags  Cap
Dev           Config     Dev  Config     Sts CSRT  (MB)
-----
DEV341_30     /dev/sd1      0030 RDF21+R-5  WD  XAM1  2063
DEV341_31     /dev/sdm      0031 RDF21+R-5  WD  XAM1  2063
DEV341_E6     N/A           00E6 RDF21+R-5  WD  XAM1  2063

```

 **NOTE:** The displayed logical device names are truncated to an existing standard length of eleven characters.

Show composite groups associated with a device group

Example

The following output shows that the **dgincg** device group is a member of two composite groups, **cg110** and **cgregular**. Also, **dgincg** is a group type of **ANY**.

To show the composite groups containing the **dgincg** device group, enter:

```

symdmg show dgincg
Group Name: dgincg
  Group Type          : ANY
  Device Group in GNS : No
  Valid               : Yes
  Symmetrix ID        : 000194900341
  Group Creation Time : Tue Dec 1 8:6:9 2009
  Vendor ID           : EMC Corp
  Application ID       : SYMCLI
  Number of STD Devices in Group : 2154
  Number of Locally-associated BCV's : 128
  Number of Locally-associated VDEV's : 0
  Number of Locally-associated TGT's : 0
  Number of Remotely-associated VDEV's (STD RDF) : 0
  Number of Remotely-associated BCV's (STD RDF) : 0
  Number of Remotely-associated TGT's (TGT RDF) : 0
  Number of Remotely-associated BCV's (BCV RDF) : 0
  Number of Remotely-associ'd RBCV's (RBCV RDF) : 0

```

```

Number of Remotely-assoc'd BCV's (Hop-2 BCV) : 0
Number of Remotely-assoc'd VDEV's (Hop-2 VDEV) : 0
Number of Remotely-assoc'd TGT's (Hop-2 TGT) : 0
Number of Composite Groups : 2
Composite Group Names : cg110
                        : cgregular

```

Perform control operations on composite group devices

Description

Use the `symcg -cg` command to perform the following operations on devices within a composite group.

Syntax

Refer to the Dell EMC Solutions Enabler CLI Reference Guide for the `symcg` command syntax and allowable control operations

GNS repository

Group name services (GNS) provides a common repository to store and maintain SYMAPI device group and composite group definitions across storage arrays that are visible to all locally attached hosts. Any host with GNS enabled can view and use the device and composite groups that are contained in the repository regardless of what host actually ran the commands to create them. This provides redundancy in case of a failure that causes the host that normally runs Solutions Enabler to lose access to the VMAX array. It also creates standards for device groups and device names that can be shared across all hosts in the environment.

Starting with Solutions Enabler V8.0.1, groups are no longer stored in a SYMAPI database file but are stored either in a local database file on the host or in a global database on the array. Both the local and global group databases are managed by the GNS daemon. The setting of SYMAPI_USE_GNS defines the group management operation mode that governs where the groups are stored:

- If SYMAPI_USE_GNS is set to DISABLE, all group information by default will be stored to a local file located in the secure directory:

```
/var/symapi/gns/storgnsd.db
```


The GNS daemon has write privileges. This file is referred to as the "local group database" or simply "local file" for groups in this document. By default, SYMAPI_USE_GNS is set to DISABLE so the group information is stored in a local database file on the host that created the group.

- If SYMAPI_USE_GNS is set to ENABLE, all group information will be stored in a global database on the array.

Enabling GNS allows group definitions to be stored on the array in a shared GNS repository. This shared GNS repository is visible to any GNS-enabled locally-attached host, enabling these hosts to perform control operations, regardless of which host initially defined the group. In addition, if one host goes down, you can still perform SYMCLI control operation from another local host in your array environment.

Solutions Enabler SYMAPI and SYMCLI do not directly access the shared repository. Instead, requests are forwarded to the GNS daemon, which processes all GNS operations. This daemon is the only entity that directly accesses the GNS shared repository and is responsible for ensuring that each host has access to the most current GNS definitions.

From each host, a GNS daemon listens for GNS requests from local clients (same host) and carries them out on the locally attached array. In addition, the GNS daemon monitors the GNS repositories on all locally-attached arrays, at a user-configured polling interval, for changes made to the shared GNS repository by other daemons (on other hosts). When a change is identified, the GNS daemon updates the host to ensure that all GNS-enabled hosts refer to the same group definitions. A set of options are available for controlling the GNS daemon. For information on configuring the GNS daemon, refer to [GNS daemon options file](#) on page 141.

 **NOTE:** If User Authorization is enabled, configuration and management of GNS requires a minimum role of StorageAdmin.

Automatic upgrade

When upgrading from an earlier version of Solutions Enabler, all active groups are automatically migrated from the existing SYMAPI database (either default or alternate) to a GNS-managed database. This migration happens once for a SYMAPI database.

Shared group definitions with GNS

A *device group* is a user-defined object comprised of devices that belong to a single array and RA groups on that device. In a GNS-enabled environment, device group definitions are stored in the GNS repository of the array on which the devices within the device group reside and are visible to all hosts locally attached to that array.

A *composite group* is also a user-defined object comprised of devices. However, the device members of a composite group can belong to multiple arrays and RA groups. In a GNS-enabled environment, composite group definitions are distributed across all arrays that contain device members of the composite group by the GNS daemon. A host must be locally attached to all arrays containing devices in the composite group to manage or control that composite group. If a host is only attached to a subset of the arrays that a composite group spans, that group will be visible to the host, but in an invalid and unusable state. This is not a recommended configuration.

The GNS state that is visible from a given host is determined by the set of arrays to which the host is connected. As seen in [Figure 5. Host-visible GNS state](#) on page 135, Host-1 is GNS-enabled and locally attached to array A, array B, and array C. Host-1 can access and modify the group definitions on all three arrays. The group definitions that are visible to Host-1 include device groups DG1, DG2, DG3, DG4, DG5, and DG6 and composite group CG1's definition, which contains devices on arrays B and C.

Host-2 is GNS-enabled and locally attached to array B and array C. The group definitions that are visible to Host-2 include device groups DG3, DG4, DG5, and DG6 and composite group CG1. Since Host-2 is locally attached to only array B and C, it cannot see any device groups or composite groups on array A (that is, DG1 and DG2).




Figure 5. Host-visible GNS state

GNS updates to group definitions

When GNS is enabled, device group and composite group definitions are stored in the common GNS repository.




Figure 6. GNS across multiple hosts

For example, in [Figure 6. GNS across multiple hosts](#) on page 136, Host-1 makes a change to composite group CG1.

1. The Host-1 GNS daemon gets this request to process.
2. The Host-1 GNS daemon updates the GNS-shared repository on all relevant arrays, in this case arrays A, B, and C.
3. The GNS daemon running on Host-2 detects the change while polling its locally-attached GNS repositories (attached arrays A, B, and C) and updates Host-2 with the updates made by Host-1.
4. Host-2's client application will detect the change on its next group call.

GNS and consistency groups

An SRDF consistency group is a composite group comprised of SRDF devices (RDF1, RDF2, or RDF21) acting in unison to preserve dependent write consistency of a database distributed across multiple SRDF systems. It maintains this consistency by using either Multi-Session Consistency (MSC) for SRDF/A or SRDF Engineuity Consistency Assist (SRDF-ECA) for SRDF/S. Both use the SRDF daemon to maintain SRDF consistency.

In a GNS-enabled environment, certain changes made to any composite group set for consistency are automatically propagated to the SRDF daemon on all relevant hosts, such as changes to the group type, SRDF group name, recovery RA group number, device

membership, device group membership, and device LdevName. When updates are made to the GNS repository, the GNS daemon updates the SRDF daemon.

For more information on setting SRDF consistency to composite groups, refer to [Create a composite group from devices in a storage group](#) on page 124.

GNS device groups and SRDF

In an SRDF scenario, a local GNS-backed device group (either RDF1 or RDF2) can be automatically mirrored on the remote array through SRDF links — for activation and use during a disaster failover situation. By default, the GNS device group definitions are stored on the local (directly attached) array — where the local devices are located.

Optionally, any changes made to the local group definition (for example, adding a standard or BCV device) can be set to automatically maintain a mirrored group definition on the remote array. This remotely mirrored group is for disaster recovery situations, where entire applications (including group aware ones), failover and are restarted on the remote side (where the remote devices are).

As changes are made to the local group definition, GNS automatically propagates corresponding changes to the remote array, so that the two are kept synchronized. For more information on configuring this option, refer to [Remote mirror device group definitions](#) on page 143.

GNS behavior in client/server mode

In client/server mode, the SYMAPI_USE_GNS setting in the SYMAPI options file on the SYMAPI server must be enabled to use GNS.

Active vs. inactive group lists

When GNS is enabled, group definitions are stored in the global GNS repository in addition to individual group database files. To facilitate the importing of groups into GNS, limited access to the group definitions held within a group database is provided while GNS is enabled.

GNS setup

Description

Use the SYMAPI_USE_GNS option in the SYMAPI options file to enable or disable GNS on each host.

Options

SYMAPI_USE_GNS

Enables or disables GNS on each host. Possible values are ENABLE or DISABLE (the default setting).

- If SYMAPI_USE_GNS is set to ENABLE, all group information is stored in a global database on each array with which specific groups are associated.
- When SYMAPI_USE_GNS is set to DISABLE, the GNS daemon stores groups in a local database file can be designated by the SYMCLI_DB_FILE environment variable. If the file is not specified, the default is used.

SYMCLI_DB_FILE

Determines the database to store groups.

The GNS daemon may store groups either in the global group database or in a local group database depending on whether SYMCLI_DB_FILE is pointing to a private database. If SYMCLI_DB_FILE is not set, then the global group database will be used to store groups. Otherwise, the local group database designated by SYMCLI_DB_FILE will be used to store groups.

NOTE: It is assumed that applications running with a private configuration database file intend to keep their modifications private, and therefore, store the groups in a private group database file.

If SYMAPI_USE_GNS is not set and if SYMCLI_DB_FILE is not pointing to a private database, the GNS daemon uses /var/symapi/gns/storgnsd.db to store groups.

When you specify an alternate database, the name of the alternate group database file is derived from the value of alternate database name and the group database file is placed under the /var/symapi/gns directory. For example, if the alternate SYMAPI database name is /usr/application/my_symapi_db.bin, the alternate group database file is /var/symapi/gns/my_symapi_db.bin_<integer>.db. The integer is a number derived from the original path name.

Examples

To enable GNS, enter:

```
SYMAPI_USE_GNS=ENABLE
```

GNS in a multihost environment

When configuring GNS in a multihost environment, it is recommended that you first review the groups currently stored on all hosts that share access to a set of arrays. Resolve any potential conflicts with other hosts. Also, if using composite groups, ensure that all hosts on which you intend to enable GNS are local to the same set of arrays to avoid composite groups appearing invalid to hosts that are not attached to all referenced arrays.

When ready, enable GNS on a single array being sure to configure any options or user authentication that you desire. For details on managing GNS through the options file, refer to the Dell EMC Solutions Enabler CLI Reference Guide. In addition to enabling GNS on the host, you can configure GNS User Authentication (set up a specific set of users with access to the GNS daemon) and GNS Daemon Control (start and stop the daemon, and configure GNS daemon options).

Once GNS is enabled on a host, activate those groups that you wish to expose using GNS. Enable GNS on a second host. Validate that the new host can see the previously activated groups and resolve any conflicts with other hosts before activating additional groups.

NOTE:

With Solutions Enabler V7.4 and higher, device and composite groups of type REGULAR are allowed to contain remote SRDF devices. However, on systems running GNS, any peer hosts running older versions of Solutions Enabler see these groups as invalid.

GNS user authentication

GNS is intended to share group definitions across multiple users and hosts in a environment via the host-installed GNS daemons. As such, access to GNS functionality is controlled by limiting permission to the GNS daemon. This access is controlled through the common daemon authorization file, `daemon_users`. This file is located in the following directories:

Table 9. Location of daemon authorization file

UNIX	<code>/var/symapi/ config/ daemon_users</code>
Windows	<code>c:\Program Files\EMC \SYMAPI\config \daemon_users</code>

Note that non-root Solutions Enabler users need to modify `c:\Program Files\EMC\SYMAPI\config\daemon_users` or `/var/symapi/config/daemon_users` to authorize the use of the GNS daemon.

NOTE:

It is important to protect this file so that only privileged administrators can modify it.

Users meeting any of the following criteria will be permitted to control and use the GNS daemon:

- User with privileges; UNIX users with root access and Windows users that are a members of the Administrators group
- Users listed in the `daemon_users` file located on each host from which they require access

Start the GNS daemon

There are three ways to start the GNS daemon:

- The daemon will be started automatically by the Solutions Enabler libraries the first time they attempt to connect with it, which can cause a slight delay in performance on that initial connection while the daemon starts and builds its cache.

NOTE: Prior to starting `storgnsd`, ensure that your default group database is current, since `storgnsd` uses the information stored in it to establish contact with your arrays.

- Start the daemon manually using the `stordaeomon` command line utility as follows:

```
stordaeomon start storgnsd
```

- Set the daemon to start automatically every time the local host is booted using the following command line:

```
stordaeomon install storgnsd -autostart
```

Prestarting the daemon, either manually or using the automatic option, is useful because the daemon may take a while to initially construct its cache — depending on the number of groups and arrays it has to load.

Restart the GNS daemon

If the GNS daemon is stopped for some reason, it can optionally be restarted automatically by an internal Solutions Enabler watchdog mechanism. A combination of the watchdog mechanism and the `-autostart` can be used to ensure that the daemon is always running, which is important with SRDF consistency composite groups to ensure that the MSC SRDF daemon is updated whenever group changes are made.

The watchdog mechanism is enabled by default. The watchdog daemon will restart the GNS daemon if it crashes or is killed. This behavior can be disabled through the `storgns:autorestart` entry in the `daemon_options` file (for more information, refer to [GNS daemon options file](#) on page 141).

The restart functionality is provided as follows:

Table 10. Restarting the GNS daemon

Unix	By a dedicated watchdog daemon (<code>storwatchd</code>) which is started automatically as needed.
Windows	By the Service Control Manager.

A maximum of three restarts within a 15-minute period will be attempted. Following a third restart (again, within a given 15-minute period), no subsequent restart will be attempted.

Manage the GNS daemon

You can control the GNS daemon with the `stordaeomon` command. The control options include:

- Stopping the daemon.
- Querying the daemon's status.
- Querying the daemon's log files.

Use the `stordaeomon -h` command to display the command syntax.

Access groups in offline mode

Description

When accessing a non-default alternate group database file (for example, from a different host), there are two ways to specify the file.

- Specify a `SYMCLI_GROUP_DB` environmental variable to indicate the path of the group database file.
- Specify a `-grpfile` command line option for `symcgs` and `syndgs` commands.

Both methods must specify a full path name of the alternate group database file. In addition, both methods are restricted to only `symcgs` and `syndg list`, `show`, `export`, and `exportall` operations.

If the specified file cannot be accessed, no error will be returned. Instead, all access attempts will return a "no groups found" error message.

Syntax

The **symcgs list, show, export** and **exportall** command syntax is as follows:

```
symcgs [-i <Interval>] [-c <Count>] [-v]
. . .
    export <CgName> [-file <FileName>] [-rdf]
           [-grpfile <GroupsDbFileName>]
exportall [-file <FileName>] [-rdf]
          [-grpfile <GroupsDbFileName>]
. . .
list [-offline] [-v]
     [-apidb | -rdf_consistency]
     [-grpfile <GroupsDbFileName>]
. . .
show <CgName> [-inactive] [-offline | -lock]
           [-grpfile <GroupsDbFileName>]
. . .
```

The **symdgs list, show, export** and **exportall** command syntax is as follows:

```
symdgs [-i <Interval>] [-c <Count>] [-v]
. . .
    export <DgName> [-delete] [-file <FileName>]
           [[-rdf [-rdfg <GrpNum>]] | [-sid <SymmID>]]
           [-grpfile <GroupsDbFileName>]
exportall [-delete] [-file <FileName>]
          [[-rdf [-rdfg <GrpNum>]] | [-sid <SymmID>]]
          [-grpfile <GroupsDbFileName>]
. . .
list [-sid <SymmId>] [-offline] [-v]
     [-grpfile <GroupsDbFileName>]
. . .
show <DgName> [-inactive] [-offline | -lock]
           [-grpfile <GroupsDbFileName>]
. . .
```

View and release GNS daemon external locks

Description

The GNS daemon uses two external locks to maintain exclusive access to the GNS repository on each array: F0 and F1.

Syntax

Use the following command to view available GNS locks:

```
symcfg -sid nnnn -lockn GNS list
```

Use the following command to manually release a GNS lock:

```
symcfg -sid nnnn -lockn GNS release
```

GNS daemon log file

Description

The GNS daemon writes its log (trace) messages to the standard location used by all daemons. The locations are:

UNIX:

/var/symapi/log/storgnsd.log0

```
/var/symapi/log/storgnsd.log1
```

Windows

```
c:\Program Files\EMC\SYMAPI\log\storgnsd.log0
```

```
c:\Program Files\EMC\SYMAPI\log\storgnsd.log1
```

These two files are written in an alternating manner. When the active one becomes full, it is closed and the other one is truncated and made active.

Syntax

Use the following syntax to display the contents of the log file:

```
stordaeomon showlog storgnsd -lines 200
```

GNS daemon options file

Description

Configuration options for the GNS daemon are contained within the daemon options file (**storgnsd**) located in the following directories:

Table 11. Locations of GNS configuration options

UNIX	/var/symapi/config/daemon_options
Windows	c:\Program Files\EMC\SYMAPI\config\daemon_options

Syntax

Use the following syntax to specify options in the file:

```
storgnsd:OptName = OptValue
```

Options

OptName

The name of the option

OptValue

The new value

Modify option file values manually

The daemon options file contains a set of parameters that can be modified to affect GNS behavior when using SYMCLI or SYMAPI commands. The file contains editable behavior parameters set to certain optional defaults in the line entries. Commented lines beginning with a pound sign (#) are ignored.

To remove any parameter option, remove the line entry, rename the file, or comment the line by adding a pound sign (#) at the beginning of the line entry.

Modify option file values via the CLI

Description

Option file values can be set using the **stordaeomon setoption** CLI utility, which programmatically makes changes to the **daemon_options** file.

By default, only options already present in the file (such as those in use or commented out) can be set. The **-force** option can be supplied to force a change, even if it requires adding a new option line to the file.

Syntax

The following is the command-line syntax for the **stordaeomon setoption** option:

```
stordaeomon setoption DaemonName  
-name OptName=OptValue [-force]
```

 **NOTE:** For information on the possible GNS daemon options file parameters, refer to the *Dell EMC Solutions Enabler CLI Reference Guide*

Options

OptName

The name of the option

OptValue

The new value


Examples

The following example shows how this utility can be used to set or change an option setting:

```
stordaeomon setoption storgnsd -name autorestart=enable
```

The following example shows how this utility can be used to remove an option setting:

```
stordaeomon setoption storgnsd -name autorestart=
```

 **NOTE:** In the above example, an empty value has been passed.

Modify option file values during runtime

Description

You can also set the daemon option value using the **stordaeomon setvar** utility, which modifies the value during runtime. The command changes the option value during runtime and will not modify the daemon options file. Upon daemon restart, the daemon option will lose the value which was set using the **setvar** command.

Syntax

The following is the syntax for the **stordaeomon setvar** option:

```
Stordaeomon setvar Daemonname -name Var=Value
```

Options

Var

The name of the option

Value

The new value

Example

For example:

```
stordaeomon setvar storgnsd -name autorestart=enable
```

Remote mirror device group definitions

The option `gns_remote_mirror` in the GNS daemon's options file determines whether GNS should attempt to remotely mirror a device group or a composite group SRDF definition contained in the shared GNS repository for a given array.

You can enable or disable the GNS remote mirror service at runtime without restarting GNS daemon. Note that the GNS daemon does not automatically pick up the change in the `daemon_options` file. Execute the `stordaeomon reload` command to make the new option setting effective.

When enabled, GNS maintains a remote mirrored group definition with the same name as the local one creating a usable group to hosts (including GNS daemons) directly connected to the remote arrays. The remote mirrored group has the same name as the local one, and has a mirror image of its contents; in other words, the data reflects the perspective of the local array. This done to ensure that the mirrored group is a legal, usable group to hosts directly connected to the remote arrays.

As a result of remotely mirroring group definitions, the following changes are made to the remote group's definition:

- The remote group definition maintains the same name and type as the local group.
- The group subtype is reversed: an RDF1 local group becomes an RDF2 remote one; an RDF2 local group becomes an RDF1 remote one.
- Attributes that are unrelated to particular devices or storage arrays are copied to the remote copy without change. For example, modification time, vendor, and group name.
- The remote RA group number (RRAGRP) attribute is changed on the remote side to contain the corresponding local SRDF group number.
- A number of attributes come in pairs: local and remote. On the remote side, these are swapped. The local and remote standard (DEVs and RDEVs), the local and remote BCV device lists (BCVs and RBCVs), and remote BCVs of the remote BCV device lists (BRBCVs and RRBCVs) are swapped.
- The device logical device name list is left unchanged. The device and remote device lists are swapped as noted above, but the logical device name list is not. The same names apply to the local device in both the original and mirrored copy.

Mirroring restrictions

Any group containing one or more of the following configurations cannot be mirrored:

- Groups containing any concurrent devices
- Groups containing any cascaded devices
- Any composite group containing a device group
- Any group containing an SRDF STD device paired with a remote BCV
- Groups containing any hop-2 devices (including 2TGT, 2BCV, 2VDEV devices)
- Device and composite groups of type=ANY
- Groups containing non-standard SRDF devices

Modify mirrored group control

An SRDF group that has been created by the GNS remote mirror service is flagged internally by GNS as being a mirror.

By default, these mirrored groups are read-only and cannot be modified or renamed—although GNS will continue to update them as the local groups upon which they are based are changed. To change this behavior, set the `SYMAPI_GNS_MIRRORED_GROUP_CONTROL` option in the SYMAPI options file to **ENABLE** (the default value is **DISABLE**).

If a mirrored group is directly modified (from a host connected to the remote array on which it is defined), the connection between it and the local group, on which it was based, is broken. At that point, it is no longer a mirror and changes to its base group (the local group) are no longer propagated to it. If a mirrored group is renamed or deleted, GNS will subsequently recreate the mirrored group from its base group.

BCV and remote BCV device alias name swap

In the remote mirror name transformation, the following changes are made:

- The BCV alias list becomes the RBCV alias list in the remote mirror. Alias names that have a value of `BCVnnn` are automatically changed to `RBCVnnn`.
- The RBCV alias list becomes the BCV alias list in the remote mirror. The alias names that have a default value of `RBCVnnn` are automatically changed to `BCVnnn`.
- The RRBCV logical list becomes the BRBCV list in the remote mirror. Logical names that have a default value of `RRBCVnnn` are automatically changed to `BRBCVnnn`.

- The BRBCV logical list becomes the RRBCV list in the remote mirror. Logical names that have a default value of BRBCV nnn are automatically changed to RRBCV nnn .

These names are the default names (logical device names) supplied by SYMCLI if you do not provide one. If you override the defaults and supply your own names that look like the above, they will still be transformed as described.

In the unlikely scenario that the you have explicitly supplied aliases, for example, BCV nnn and RBCV nnn , it is possible that the above transformation may result in duplicate alias names in the remote mirror. For example, if the original BCV alias list contained **BCV001, RBCV001**, the remote mirror's RBCV alias list would contain **RBCV001, RBCV001**, which would cause a name conflict.

The following example illustrates the mirrored group data:

```

--- Local DG Group, on Symm1 ---
Type                               RDF1
Role                               Have a Remote Mirror
Remote Symm                         Sym2
Local RDFGrp                         2
Remote RDFGrp                       5
Devices                             0001,0002,0003
Remote-Devs                         0091,0092,0093
LDevNames                           DEV001,DEV002,DEV003

BCVs                                0010,0011
Remote-BCVs                         0020,0021,0022
BCV LdevName                        BCV001,BCV002
RBCV LdevName                        RBCV001,RBCV002,RBCV003

Vendor                              EMC
ModTime                             12345
GateKeepers                         /dev/rdsk/c2t0d2

```

```

--- Remote DG Group, on Symm2 ---
Type                               RDF2
Role                               Am a Remote Mirror
Remote Symm                         Sym1
Local raGrp                         5
Remote raGrp                       2
Devices                             0091,0092,0093
Remote-Devs                         0001,0002,0003
Dev Aliases                         DEV001,DEV002,DEV003

BCVs                                0020,0021,0022
Remote-BCVs                         0010,0011
BCV Aliases                         BCV001,BCV002,BCV003
RBCV Aliases                        RBCV001,RBCV002

Vendor                              EMC
ModTime                             12345

```

Identify GNS groups

After configuring your array environment to support the GNS option, you can begin to create and modify groups. Use **syndg show DgName**, for device groups, or **symcg show CgName**, for composite groups, to determine if a group is stored in the GNS group list.

Identify device group definition

Description

For a device group, the output of the **syndg show DgName** command identifies whether a group definition is stored in GNS.

Example

```
syndg show MyRegDeviceGroup
```


Sample output

This sample output shows that group **MyRegDeviceGroup** is stored in the GNS list since the output parameter **Device Group in GNS** has a value of **Yes**. It also shows the device group's GNS mirror state.

```
Group Name: MyRegDeviceGroup

Group Type                : REGULAR
Device Group in GNS      : Yes (Is Mirror)
Valid                    : Yes
Symmetrix ID             : 000184500160
Group Creation Time      : Tue Aug 4 16:44:18 2006
Vendor ID                : EMC Corp
Application ID           : SYMCLI

Number of STD Devices in Group : 20
Number of Locally-associated BCVs : 20
Number of Locally-associated VDEVs : 0
. . .
```

Identify composite group definition

Description

For a composite group, the output of the **symcgs show CgName** command identifies whether a group definition is stored in GNS.

Example

The sample output below shows that group **MyCompGroup** is stored in the GNS list since the output parameter **CG in GNS** has a value of **Yes**. It also shows the composite group's GNS mirror state. The group definition has been enabled for SRDF consistency using MSC.

```
% symcgs show MyCompGroup

Composite Group Name: MyCompGroup

Composite Group Type      : RDF1
Valid                    : Yes
CG in PowerPath          : No
CG in GNS                 : Yes (Is Mirror)
RDF Consistency Protection Allowed : Yes
RDF Consistency Enabled  : Yes

Number of RDF (RA) Groups : 4
Number of STD Devices     : 64
Number of BCVs (Locally-associated) : 0
Number of VDEVs (Locally-associated) : 0
Number of RBCVs (Remotely-associated STD-RDF) : 0
Number of BRBCVs (Remotely-associated BCV-RDF) : 0
Number of RRBCVs (Remotely-associated RBCV) : 0
. . .
```

GNS backup and restore mechanism

You can set the GNS daemon to automatically make a backup copy of local or global groups in a backup file. For global groups, use the backup copy to restore the global database on arrays. For local groups, use the backup copy to restore the groups when the primary group database file is corrupted. This backup file is *N* hours behind the in-memory database where *N* is a configurable backup interval with a default value of 6 hours.

Use the following daemon options to control the backup process:

To enable/disable backups

Set **GNS_DB_BACKUP** to either **DISABLE** or **ENABLE** where **DISABLE** is the default. When you set this option to **ENABLE**, the GNS daemon automatically backs up both the default local database and the global database into a backup file.

To set the backup interval time

Set the `GNS_DB_BACKUP_INTERVAL` to an integer value in hours. This option is only meaningful if `GNS_DB_BACKUP` option is set to `ENABLE`. The default backup interval is 6 hours. The valid range for the interval is (1..720).

All GNS database backup files are stored in the `/var/symapi/gns` directory. For both local and global GNS databases, up to two copies of backup files are maintained:

- For the GNS local database, the backup file names are `local_group_db.backup` and `local_group_db.backup.bak`, respectively.
- For the GNS global database, the backup file names are `global_group_db.backup` and `global_group_db.backup.bak` respectively.

Manual backup to the GNS database

Description

To manually backup to the GNS database, use the `stordaemon` program.

Syntax

Use the following syntax for local backup:

```
stordaemon action storgnsd -cmd backup_db local
```

Use the following syntax for global backup:

```
stordaemon action storgnsd -cmd backup_db global
```

Manual restore from the GNS database

Description

To manually restore from the GNS database, use the `stordaemon` program.

Syntax

Use the following syntax for a local restore:

```
stordaemon action storgnsd -cmd restore_db local
```

Use the following syntax for a global restore:

```
stordaemon action storgnsd -cmd restore_db global changed_records:sid  
stordaemon action storgnsd -cmd restore_db global all_records:sid
```

The `changed_records` option only updates those records that differ from group records in the backup file. The `all_records` option rebuilds the entire database using the backup file as source data.

Troubleshoot GNS

When using GNS, be aware of the possibility of group name collision and invalid groups and what SYMAPI does to rectify each scenario.

Group name collisions

Each GNS daemon attempts to ensure that group names are unique (relative to a group type). However, duplicate-named groups may occur.

If users at two different hosts simultaneously create groups of the same type with the same name using devices on different arrays, both attempts might succeed. Because each host's GNS daemon is polling for changes made elsewhere, the daemons might not detect the name conflict until after the fact. A subsequent consideration change could then allow the GNS daemon to see both groups.

 **NOTE:** For information on renaming composite groups, see [Rename a composite group](#) on page 130.

Example

If a GNS daemon detects duplicate group names, it returns a state flag informing clients of this fact and modifies the group names to be unique by attaching a numeric suffix. Duplicate device groups with the name **MyDG** could be modified in the following way:

```
MyDG#12345678
MyDG#87654321
```

While in this state (duplicate names), the only operations permitted are those used to resolve the name conflict: group rename and delete (with the **-force** flag). For example:

```
syndg delete MyDG#12345678 -force
syndg rename MyDG#87654321 MyDG
```

Invalid groups

A composite group is marked as invalid if one or more of the arrays that it resides on (in a GNS-enabled environment) cannot be reached. For example, each array that a consistency group spans records the identity of the other arrays that the group spans. If a discrepancy is noticed where a composite group is not defined on an array, even though some other array indicates it should be there, that group is assumed to be invalid.

Groups are also placed into an invalid state if certain internal bookkeeping information maintained by GNS is found to be missing or incorrect. For example, a consistency group spans two arrays and the GNS state recorded on each array verifies that the group also resides on the other array. If at some point one array is found to not contain the group (while the other array still thinks it should be found there), the group is assumed to be invalid.

This could happen if one of the arrays is reinitialized or a host attached only to one array decides to delete the composite group there. Since the group definition is also stored on an unreachable array, that group state is seen as invalid from that host. A user there would have to use the **-force** flag to delete the group. Groups marked as invalid can be deleted with the **-force** flag.

List GNS data about groups

The GNS **stordaeon list** command supports the following features:

- **stordaeon action storgnsd -cmd list_groups**
Lists groups based on GNS daemon mode. If the GNS daemon is running with USE_GNS enabled, it will list all global groups. Otherwise, it will list all local groups in **storgnsd.db**.
- **stordaeon action storgnsd -cmd list_groups global**
Lists all global groups.
- **stordaeon action storgnsd -cmd list_groups local**
Lists all local groups.
- **stordaeon action storgnsd -cmd list_groups local:SYMCLI_DB_FILE**
Lists groups in a private GNS database determined by the environment variable SYMCLI_DB_FILE.

Show GNS data about groups

The GNS **stordaeon show** command supports the following features:

- **stordaeon action storgnsd -cmd show_group cg: name**
Shows composite groups based on GNS daemon mode. If the GNS daemon is running with USE_GNS enabled, it will list all global groups. Otherwise, it will list all local groups in **storgnsd.db**.
- **stordaeon action storgnsd -cmd show_group dg: name**
Shows device groups based on GNS daemon mode. If the GNS daemon is running with USE_GNS enabled, it will list all global groups. Otherwise, it will list all local groups in **storgnsd.db**.
- **stordaeon action storgnsd -cmd show_group cg: name:local**
Shows all local composite groups.

- **stordaemon action storgnsd -cmd show_group dg: name:local**
Shows all local device groups.
- **stordaemon action storgnsd -cmd show_group cg: name:global**
Shows all global composite groups.
- **stordaemon action storgnsd -cmd show_group dg: name:global**
Shows all global device groups.
- **stordaemon action storgnsd -cmd show_group cg: name:local:SYMCLI_DB_FILE**
Shows composite groups in a private GNS database determined by the environment variable SYMCLI_DB_FILE.
- **stordaemon action storgnsd -cmd show_group dg: name:local:SYMCLI_DB_FILE**
Shows device groups in a private GNS database determined by the environment variable SYMCLI_DB_FILE.

Inline Compression

This chapter describes how to enable or disable compression on a storage group, and how to monitor the compression ratio on a storage group or at the SRP level.

Topics:

- [Inline compression overview](#)
- [Inline compression control operations](#)
- [Inline compression display and reporting](#)

Inline compression overview

For VMAX All Flash arrays running HYPERMAX OS 5977 Q3 2016 or higher, user data can be compressed. Solutions Enabler enables and disables compression on a storage group, and, when enabled, monitors the current Data Reduction Ratio (DRR) on the storage group or at the SRP level. Compression is supported on open systems (FBA) only, including eNAS.

NOTE: All Flash arrays require compression hardware and proper sizing to accommodate upgrading to and enabling compression. The compression attribute is enabled by default for all FAST managed storage groups, with all FBA devices, if the SRP for the storage group supports compression.

Some key features of compression are:

- When enabling compression on a storage group, all incoming writes are considered for compression.
- When disabling compression on a storage group, all incoming writes stop the compression of new data.
- If an All Flash array is upgraded to HYPERMAX OS 5977 Q3 2016 software, existing storage groups can be enabled for compression and incoming data writes are considered for compression.

NOTE: Upgraded All Flash arrays require a conversion script (run by customer service personnel) to enable compression.

- All data services are supported (such as, SnapVX, SRDF, D@RE).

Compression is not supported on VMAX3 hybrid arrays (100K, 200K, 400K) or External Flash arrays (no FAST.X support).

Inline compression control operations

On VMAX All Flash arrays with compression feature support, compression is enabled by default.

The `symmsg` CLI command controls the following compression operations:

- Removes compression when creating a storage group using the `symmsg create -nocompression` command.
- Sets or removes compression on an existing storage group using the `symmsg set` command.
- Exports the compression attribute (`P <compression enabled>`) on the storage group to a text file on the host when using the `symmsg export/exportall` commands.
- Imports the compression attribute on FAST managed storage groups when using the `symmsg import/importall` commands.

When importing a storage group and the target is a VMAX Array running HYPERMAX OS 5977 Q3 2016SR and higher:

- If the imported storage group already exists on the array, the command fails with the error message "Cannot use the specified name because it's already in use".
- If the SRP does not exist on the target array, the storage group is imported and set to use the Optimized service level with no Workload. The compression attribute is cleared if the SRP is not enabled for compression.
- If the service level set on the storage group does not exist on the target array or cannot be supported based on the SRP that is currently used, the storage group is imported and set to use the Optimized service level with no Workload. The compression attribute is cleared if the SRP is not enabled for compression.
- If the both the service level and the SRP set on the storage group do not exist on the target array, the storage group is restored with no service level or SRP and is not FAST managed. The compression attribute is cleared.

If the target is a VMAX array running HYPERMAX OS lower than 5977 Q3 2016SR, the service level, Workload, SRP names, and compression attribute are copied, but the compression attribute is cleared.

Set or remove compression at storage group level

Description

When creating a storage group (`symmsg create`), the compression attribute is enabled by default on FAST managed storage groups if the associated SRP supports compression. For storage group to be FAST managed either the `-sl` or `-srp` option must be specified, otherwise compression is disabled.

For existing storage groups, the compression attribute can be enabled or disabled with the `symmsg set` command using the `-compression` or `-nocompression` options. An error is returned with the `-compression` option if the associated SRP does not support compression.

NOTE: Setting or removing compression on storage groups requires the following permissions:

- Access Type – Base
- Authorization Rights – Storage Admin

Syntax - symmsg create

To override the default compression setting when creating a storage group, use the following syntax:

```
symmsg -sid <SymmID> [-i <Interval>] [-c <Count>] [-v]
  create <SgName>
    [-bw_max <MBperSec>]
    [-iops_max <IOperSec>]
    [-dynamic <NEVER | ALWAYS | ONFAILURE>]
    [[-sl <SLName> [-wl <WorkloadName>]]]
    [-srp <SRPName>] [-nocompression]
```

The `-nocompression` option can be abbreviated as `-noc`.

Refer to [Create storage groups](#) on page 95 for complete command syntax and options.

Syntax - symmsg set

To set or remove compression on an existing storage group use the following syntax:

NOTE: When setting the compression option the storage group must be FAST managed or the command will fail.

```
symmsg -sg <SgName> -sid <SymmID> [-i <Interval>]
  [-c <Count>]

  set [<[-bw_max <MBperSec> | NOLIMIT ]
    [-iops_max <IOperSec> | NOLIMIT ]
    [-dynamic <NEVER | ALWAYS | ONFAILURE>]>]
    [-sl <SLName> [-wl <WorkloadName>] | -nosl]
    [-srp <SRPName> | -nosrp]
    [-compression | -nocompression]
```

NOTE: The `-compression` option can be abbreviated as `-com` and `-nocompression` as `-noc`.

Refer to [Modify storage group properties](#) on page 100 for complete command syntax and options.

Examples

To set compression on storage group `Test_SG_1`, enter:

```
symmsg set -sg Test_SG_1 -sid 146 -com
```

To set remove compression on storage group `Test_SG_1`, enter:

```
symmsg set -sg Test_SG_1 -sid 146 -noc
```

Remove compression at storage container resource level

Description

The default compression attribute can be removed using the `-nocompression` option.

NOTE: The compression attribute is cleared if the associated SRP does not support compression.

Syntax

To remove compression for a resource when adding it a storage container, use the following syntax:

```
symcfg -sid <SymmID> -sc -sc_name <StorageContainer>
      [-noprompt]

      add -sresource <StorageResourceName>
          <-sl <SLName> [-wl <WorkloadName>] [-nocompression]>
          [-srp <SRPName>]
          <-subscribed_max <GB>>
```

The `-nocompression` option can be abbreviated as `-noc`.

Refer to [Add storage resource to storage container for VVols](#) on page 240 for complete command syntax and options.

Inline compression display and reporting

Compression display

The following commands display the compression attribute on the storage group and storage container resource:

- `symmsg list` – Refer to [List storage groups](#) on page 110 for output display.
- `symcfg list -sc` – Refer to [List storage containers for VVols](#) on page 244 for output display.
- `symcfg show -sc` – Refer to [Show storage container for vVols](#) on page 243 for output display.
- `symmsg show` – Refer to [View storage group details](#) on page 113 for output display.

Compression reporting

- `symcfg list -tdev` and `symcfg list -tdev -detail` – report the compression ratio (`Comp Ratio`) of thin devices. The reported device compression ratio is the ratio of Logical Allocated Capacity to the Physical Allocated Capacity. Logical Allocated Capacity is the total allocated capacity calculated using the device logical track size, and the Physical Allocated Capacity is the total allocated capacity calculated using the compressed pool track size. Refer to [View thin devices](#) on page 77 or [View thin device details](#) on page 78 for output display.
- `symmsg show` – reports the compression attribute (`Compression Enabled`) and the compression ratio (`Compression Ratio`) for storage groups. The reported compression ratio is the ratio of Logical Allocated Capacity of all the devices in the storage group to the Physical Allocated Capacity of all devices in the storage group. Logical Allocated Capacity is the total allocated capacity calculated using the device logical track size, and the Physical Allocated Capacity is the total allocated capacity calculated using the compressed pool track size. Refer to [View storage group details](#) on page 113 for output display.
- `symcfg list -pool` – reports the compression ratio (`Comp Ratio`) for each of the thin pools in the array and the total compression ratio for the entire array. The ratio on Total line represents ratio for whole system based on current pools utilization. The calculation is done based on used tracks in each pool. Since DSE tracks are never compressed they are not part of ratio calculation. The Compression Ratio in this case is ratio between sum of Logical Used Pool Capacity and sum of Physical Used Pool Capacity. The Logical Used Pool Capacity is number of used tracks in the pool not counting DSE tracks multiplied by logical track size of 128K. The Physical Used Pool Capacity is number of used tracks not counting DSE tracks multiplied by pool track size. Refer to [View all thin device pools](#) on page 75 for output display.
- `symcfg show -thin -pool -detail` – reports the compression state (`Compression State`) and the compression ratio (`Compression Ratio`) for a specified pool. For HYPERMAX OS 5977 and above, the pool ratio is constant for a given pool and calculated as the ratio of the device logical track size track to track size of this pool. Refer to [View thin device pool details](#) on page 76 for output display.
- `symcfg list -srp` and `symcfg show -srp` – report the compression state (`Compression State`) and the Data Reduction ratio (`Data Reduction Ratio`) for SRPs in the array. The reported DRR ratio is the ratio of Logical Allocated Capacity of all

devices in SRP and the Physical Allocated Capacity for all devices in SRP. Logical Allocated Capacity is the total allocated capacity calculated based on the device logical track size and the Physical Allocated Capacity is the total allocated capacity calculated based on the compressed pool track size. Refer to [List Storage Resource Pools details](#) on page 163 for output display.

- Efficiency reports– report the overall system efficiency, and the compression efficiency of the thin devices in the SRP. Refer to [Report system efficiency](#) on page 154 and [Report SRP inline compression efficiency](#) on page 154 for syntax, example code and output display.
- `symcfg show -compression_conversion_status` – reports compression conversion start time, status, and percent complete. Refer to [Report array compression conversion status](#) on page 155 for syntax, example code and output display.

Report data compressibility

Description

The `symcfg list -sg_compression` reports the maximum data compressibility on a storage group for compression capable arrays. The reported compressibility is an estimate calculated over the previous 24 hours. The report of storage groups is sorted by the storage group name by default. There is an option to report the compressibility of storage groups in descending order. By default, storage groups already enabled for compression are not reported, but there is an option to list both compression capable storage groups and compression enabled storage groups.

Options

`-by_compressibility`

Sorts the report by compressibility in descending order. The `-by_com` abbreviated is allowed.

`-all`

Reports both compression capable storage groups and compression enabled storage groups.

`-srp`

Filters report for storage groups associated only with specified the Storage Resource Pool (SRP). For devices associated with the SRP but not in any storage group these devices are reported as `not_in_sg`.

Examples

To list the compressibility for the storage groups on array 188, enter:

```
symcfg list -sid 188 -sg_compression
```

Sample output

```
S T O R A G E   G R O U P S
```

```
Symmetrix ID:          000197800188
```

```
Name                   : SRP_1
```

Storage Group Name	Number Devices	Allocated (GB)	Used (GB)	Estimated Ratio
Customer	106	350.0	350.0	1.2:1
<not_in_sg>	54	140.5	140.5	3.4:1

```
Name                   : SRP_2
```

Storage Group Name	Number Devices	Allocated (GB)	Used (GB)	Estimated Ratio
Payroll	200	756.3	756.3	2.2:1
<not_in_sg>	54	140.5	140.5	3.4:1

Output with `-all` option:

```
S T O R A G E   G R O U P S
```



```

Symmetrix ID:          000197800188

Name                   : SRP_1

Storage Group Name     Flags  Number  Allocated    Used  Estimated
                       C      Devices (GB)         (GB)  Ratio
-----
AcctPay                X           6      40.0        20.5   5.2:1
Customer               .          106     350.0       350.0   1.2:1
<not_in_sg>           .           54     140.5       140.5   3.4:1

Name                   : SRP_2

Storage Group Name     Flags  Number  Allocated    Used  Estimated
                       C      Devices (GB)         (GB)  Ratio
-----
Payroll                .          200     756.3       756.3   2.2:1
<not_in_sg>           .           54     140.5       140.5   3.4:1

Legend:
Flags:
(C)ompression          X = Compression Enabled, . = N/A

```

Efficiency reporting

The following efficiency ratio and percent saved reporting for the SRP and the overall system is available:

- Virtual Provisioning efficiency
- Snapshot efficiency
- Compression efficiency
- Overall efficiency

Efficiency ratios are reported in units of 1/10th:1

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

Reported values

For both the overall system efficiency and SRP efficiency reporting the following values are reported:

Virtual Provisioning Efficiency:

- Saved (%) – Percentage savings of the TDEV configured storage presented to the hosts and the TDEV allocated storage.
- Shared Ratio – Ratio of the TDEV allocated storage and the TDEV Logical Backend Storage.
- VP Overall Efficiency Ratio – Ratio of the TDEV configured storage and the TDEV Logical Backend Storage.

Snapshot Efficiency

- Saved (%) – Percentage savings of the sum of all TDEV snapshot sizes (at the time of snapshot creation) and the TDEV Snapshot Allocated Space.
- Shared Ratio – Ratio of the Snapshot Allocated Storage and the RDP Logical Backend Storage.
- Snapshot Overall Efficiency Ratio – Ratio of the sum of all Snapshot sizes and the RDP Logical Backend Storage.

Compression Efficiency

- Compression VP Ratio (%) – Ratio of the TDEV Logical Backend Storage and the TDEV Physical Used Storage.
- Compression Snapshot Ratio – Ratio of the RDP Logical Backend Storage and the RDP Physical Used Storage of the RDP space.
- Compression Overall Ratio – Ratio of the sum of all TDEVs + RDP Logical Backend Storage and the TDEVs + RDP Physical Used Storage.

Overall Efficiency

- Overall Efficiency Ratio – Ratio of the sum of all TDEVs + Snapshot sizes and the Physical Used Storage.

Report system efficiency

Description

The `symcfg list` command, when used with the `-efficiency` option, reports system efficiency.

Syntax

To list system efficiency using the `symcfg list` command, use the following syntax:

```
symcfg [-sid <SymmID>] [-offline] list -efficiency
```

NOTE: The `-efficiency` option can be abbreviated as `-eff`.

Examples

To list system efficiency for array 084 using the `symcfg list` command, enter:

```
symcfg list -sid 084 -efficiency
```

Sample output

For `symcfg list -efficiency` command:

```
SYMMETRIX EFFICIENCY
-----
SymmID          Overall  Data Reduction      VP Snapshot
                Ratio    Ratio Enabled      Ratio    Ratio
-----
000197800188    8.4:1   2.5:1      100    10.0:1   1.4:1
```

Report SRP inline compression efficiency

Description

The `symcfg list` command, when used with the `-efficiency` and `-srp` options, reports SRP efficiency.

Syntax

To list SRP efficiency using the `symcfg list` command, use the following syntax:

```
symcfg [-sid <SymmID>] [-offline] list -srp -efficiency
```

NOTE: The abbreviation `-eff` for `-efficiency` is allowed

Examples

To list SRP efficiency for array 099 using the `symcfg list` command, enter:

```
symcfg list -sid 099 -srp -efficiency
```

Sample output

For `symcfg list -srp -efficiency` command:

```
STORAGE RESOURCE POOLS
Symmetrix ID : 000197800188
-----
Name          Overall  Data Reduction      VP Snapshot
                Ratio    Ratio Enabled      Ratio    Ratio
```

```
-----  
SRP_1                8.4:1   2.5:1   100   10.0:1  1.4:1  
-----
```

Report array compression conversion status


Description

The compression conversion status report requires the `-sid` option. This information is not stored in the Solutions Enabler database so it is not available for offline access.

Syntax

To report the compression conversion status for an array, use the following syntax:

```
symcfg [-sid <SymmID>] -compression_conversion_status
```

 **NOTE:** The abbreviation `-eff` for `-efficiency` is allowed

Examples

To report the compression conversion status for array 456, enter:

```
symcfg show -compression_conversion_status -sid 456
```

Sample output

```
Symm.....: 1971123456  
Conversion Start.....: 06/14/2016 14:55:12  
Status.....: Active  
Percent Complete.....: 93.2%
```

Fully Automated Storage Tiering

This chapter describes Fully Automated Storage Tiering (FAST).

Topics:

- [Fully Automated Storage Tiering](#)
- [Storage Resource Pool management](#)
- [FAST information reporting](#)

Fully Automated Storage Tiering

NOTE:

This section describes FAST operations for VMAX3 arrays. The EMC VMAX3 Family Product Guide for VMAX 100K, VMAX 200K, VMAX 400K with HYPERMAX OS provides additional details about FAST concepts and operations.

EMC Fully Automated Storage Tiering (FAST) provides automated management of VMAX 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.

FAST technology moves the most active parts of 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 and helps reduce acquisition, power, cooling, and footprint costs. FAST factors 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 provides the ability to deliver variable performance levels through *Service Levels*. Thin devices are added to *Storage Groups* and the storage group are assigned a specific Service Level which sets performance expectations.

A Service Level is the response time target for a storage group. The Service Level sets the VMAX 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.

There are five Service Levels. For each one, except Optimized, the workload type, such as Online Transaction Processing (OLTP) or Decision Support System (DSS), is specified. The Optimized service level does not allow for setting the I/O type. A Service Level cannot be modified, however a storage group can be assigned based on the required service level.

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

VMAX3 arrays running HYPERMAX OS 5977 are custom-built and pre-configured with array-based software applications, including a factory pre-configuration for FAST that consists of:

- Data device (TDAT) — an internal device that is dedicated to providing physical storage used by thin devices.
- Data pool — a collection of data devices of identical emulation and protection type, all of which reside on disks of the same technology type and speed. The disks in a data pool are from the same disk group.
- Disk group — a collection of physical drives within the array that share the same performance characteristics.
- Storage Resource Pool (SRP) — one (default) FAST SRP is pre-configured on the array. This process is automatic and requires no setup.

Storage Resource Pool management

The `symconfigure set` command modifies Storage Resource Pool configuration. Configuration modifications include:

- set reserve capacity
- enable or disable SRDF/A DSE
- change SRP descripton

Modify storage resource pools

Syntax

To modify a Storage Resource Pool(SRP), use the following syntax:

```
symconfigure set srp <SRP Name>,  
<[resv_cap = <n | NONE>]  
[,rdfa_dse = <ENABLE | DISABLE]>  
[,description = 'SRP Description']>;
```

The following security privileges are required to execute the `symconfigure set` command:

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin.

Options

resv_cap

A percentage of the capacity of the SRP reserved for device 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 SRP 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 device write I/O activities.

NOTE: Existing TimeFinder snapshot sessions created on devices in the SRP are invalid if the free capacity of the SRP, as a percentage of the usable capacity, goes below the reserved capacity.

rdfa_dse

There is always one SRP assigned available for SRDF/A DSE allocations. Valid values are enable and disable. The default SRP for FBA emulation is used by default. Enabling `rdfa_dse` on a SRP will make that pool available for SRDF/A DSE allocations (and implicitly sets the currently assigned SRP to "disabled"). The maximum amount of storage from a SRP allowed for DSE is controlled by the system wide `dse_max_cap` setting, as described in the Dell EMC Solutions Enabler SRDF Family CLI User Guide

NOTE: This option is not allowed for SRPs with only external storage.

description

SRP description. Maximum description length is 127 characters (excluding null termination character). Valid values are "a - z", "A - Z", "0 - 9", underscore, hyphen "-", space, period ".", and comma ",".

Examples

To set the existing SRP reserve capacity to a value of 10% and enable SRDF/A DSE operations, enter:

```
symconfigure -sid 230 commit -cmd "set srp Primary_SRP,  
resv_cap=10, rdfa_dse=ENABLE;"
```

To modify the description of an SRP, enter:

```
symconfigure -sid 087 commit -cmd "set srp DEFAULT_SRP description = 'SRP with description  
modification';"
```

```
Execute a symconfigure operation for symmetrix '000197100087' (y/[n]) ? y  
A Configuration Change operation is in progress. Please wait...  
Establishing a configuration change session.....Established.  
Processing symmetrix 000197100087  
{  
set srp DEFAULT_SRP description = 'SRP with description modification';  
}  
Performing Access checks.....Allowed.  
. . .
```

```
Terminating the configuration change session.....Done.  
The configuration change session has successfully completed.
```

Restrictions:

- New parameters must be different than existing parameters.
- The target SRP for the operation cannot be the default SRP for FBA emulation, and the specified state cannot be "disabled".

Rename storage resource pools

Syntax

To rename a Storage Resource Pool (SRP), use the following syntax.

```
symconfigure rename srp <SRP Name> to <New SRP Name>;
```

The following security privileges are required to execute the `symconfigure rename` command:

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin.

Options

New SRP Name

Modified SRP name. Maximum name length is 32 characters. Valid characters are only alphanumeric characters, hyphens "-", and underscores "_", with leading hyphen or underscore not allowed.

Example

```
symconfigure -sid 087 commit -cmd "rename srp DEFAULT_SRP to SRP_RENAMED;"
```

```
Execute a symconfigure operation for symmetrix '000197100087' (y/[n]) ? y  
A Configuration Change operation is in progress. Please wait...  
  Establishing a configuration change session.....Established.  
  Processing symmetrix 000197100087  
  {  
rename srp DEFAULT_SRP to SRP_RENAMED;  
  }  
  Performing Access checks.....Allowed.  
  .  
  .  
  Terminating the configuration change session.....Done.  
The configuration change session has successfully completed.
```

Restrictions

In a single command file that includes "set srp" and "rename srp" commands, the following rules apply:

- Allows using the old SRP name for "set srp" commands and renaming the SRP as the last command.
- Allows renaming the SRP as the first command and using the new name for the subsequent "set srp" commands.
- Does not allow using the old SRP name for some commands, changing the SRP name, and then using the new SRP name for subsequent commands.

FAST information reporting

FAST reporting lists the Service Level, storage groups, and Storage Resource Pools configured on the array, and also generates reports detailing the demand that storage groups or Service Level are placing on the Storage Resource Pools.

List Service Level

Description

The `symcfg list` and `symcfg show` commands report Service Levels and available workloads.

Options

- all**
Displays only the Service Levels that are allowed for the specific VMAX array.
- v**
Provides a verbose list of Service Levels and workloads with additional descriptions.
- fba**
Lists only the Service Levels with FBA emulation.
- ckd**
Lists only the Service Levels with CKD emulation. This option available with HYPERMAX OS 5977 Q12106SR or higher.

 **NOTE: CKD devices support only Diamond, Bronze, and Optimized Service Level.**

Example

To list Service Level details on array 063, enter:

```
symcfg -sid 063 list -sl -detail -all
```

Sample Output

Output with `-all` filter:

```

                SERVICE LEVEL
Symmetrix ID      : 000197200063

```

Name	Workload	Approx Resp Time (ms)	Service Level Base Name
Optimized	N/A	N/A	Optimized
Diamond	OLTP	0.8	Diamond
Diamond	OLTP_REP	2.3	Diamond
Diamond	DSS	2.3	Diamond
Diamond	DSS_REP	3.7	Diamond
Diamond	<none>	0.8	Diamond
Platinum	OLTP	3.0	Platinum
Platinum	OLTP_REP	4.4	Platinum
Platinum	DSS	4.4	Platinum
Platinum	DSS_REP	5.9	Platinum
Platinum	<none>	3.0	Platinum
Gold	OLTP	5.0	Gold
Gold	OLTP_REP	6.5	Gold
Gold	DSS	6.5	Gold
Gold	DSS_REP	7.9	Gold
Gold	<none>	5.0	Gold
Silver	OLTP	8.0	Silver
Silver	OLTP_REP	9.5	Silver
Silver	DSS	9.5	Silver
Silver	DSS_REP	10.9	Silver
Silver	<none>	8.0	Silver
Bronze	OLTP	14.0	Bronze
Bronze	OLTP_REP	15.5	Bronze
Bronze	DSS	15.5	Bronze

Bronze	DSS_REP	16.9	Bronze
Bronze	<none>	14.0	Bronze

Show specific Service Level

Description

The `symcfg show -sl` command displays a specified Service Level.

Options

- fba**
Shows only FBA devices for specified Service Level.
- ckd**
Shows only CKD devices for specified Service Level.

Example

To show information about the Platinum Service Level for array 063, enter:

```
symcfg -sid 063 show -sl Platinum
```

Sample Output

```
Symmetrix ID      : 000197200063

Name              : Platinum
Service Level Base Name : Platinum

Workloads (5)
{
    Name              Approx
                   Resp
                   Time
                   (ms)
    -----
    OLTP              3.0
    OLTP_REP          4.4
    DSS               4.4
    DSS_REP           5.9
    <none>            3.0
}
```

List storage groups

NOTE:

Storage groups on page 95 explains how to create storage groups and how to add and remove devices from a group.

Description

The `symmsg list -detail` command shows details for all storage groups, including Service Level, workload, and Storage Resource Pool (SRP) configured for each storage group.

Options

- by sl**
Sorts storage group information by Service Level.
- by srp**

Sorts storage group information by SRP.

Example

To list storage group details for array 063, enter:

```
symmsg -sid 063 list -detail
```

Sample output

NOTE: Starting with Solutions Enabler 8.0.1 the **FAST R** column for listing RDF coordination (enabled or disabled) has been removed from the output.

- If the Service Level or SRP name length exceeds 24 characters output display truncates the name to 23 characters and displays "*" for the 24th character.
- When the (F)ast flag is set, it indicates that the storage group has a Service Level or a SRP set and it is FAST-managed.
- A FAST-managed storage group that reports <none> for the Service Level Name uses the Optimized Service Level.
- A FAST-managed storage group that reports <none> for the SRP Name uses the array's default Storage Resource Pool for the emulation.

S T O R A G E G R O U P S

Symmetrix ID: 000197200063

Storage Group Name	Flags	Number	Number	Child	Service Level Name	Workload	SRP Name
	EFM SLC	Devices	GKs	SGs			
151SG	F.X ...	10	10	0	<none>	<none>	<none>
EMBEDDED_NAS_DM_SG	FXX ...	8	2	0	Gold	<none>	<none>
Enasvt	F.X ...	5	0	0	<none>	<none>	<none>
ESG253_SG	FXX ...	6	6	0	Optimized	<none>	SRP_1
ESX253_SG	F.X ...	25	24	0	<none>	<none>	<none>
SVT_ISCSI_SG	F.X ...	3	1	0	<none>	<none>	<none>
V2SG1	FXX ...	1	0	0	Gold	<none>	<none>
V2SG_LINUX	FXX ...	1	0	0	Gold	<none>	<none>
V2SG_SOL	FXX ...	1	0	0	Gold	<none>	<none>
V2SG_U4V	FXX ...	1	0	0	Gold	<none>	<none>

Legend:

Flags:

Device (E)mulation A = AS400, F = FBA, 8 = CKD3380,
 9 = CKD3390, M = Mixed, . = N/A
(F)ast X = Fast Managed, . = N/A
(M)asking View X = Contained in Mask View(s), . = N/A
Cascade (S)tatus P = Parent SG, C = Child SG, . = N/A
Host IO (L)imit D = Defined, S = Shared, B = Both, . = N/A
(C)ompression X = Compression Enabled, . = N/A

List Storage Resource Pools

Options

-fba

Lists only the SRPs with FBA emulation.

-ckd

Lists only the SRPs with CKD emulation. This option is available for arrays running HYPERMAX OS 5977 Q12016SR or higher.

Examples

To list the details for all SRPs configured on array 087, enter:

```
symcfg -sid 087 list -srp -detail
```

To list the details for all SRPs configured on array 087 with -fba emulation, enter:

```
symcfg -sid 087 list -srp -fba
```

Sample output

Output with **-detail** option (reports capacity metrics) :

```

      STORAGE RESOURCE POOLS

Symmetrix ID   : 000197100087

                                     C A P A C I T Y
-----
Name              Flg      Usable  Allocated   Free   Subscribed
                   DR      (GB)    (GB)        (GB)    (GB)
-----
DEFAULT_SRP      FX       1847.4   515.1     1332.3   80065.6
TEST_2           CX       2630.1    0.0     2630.1    0.0
-----
Total              4477.5   515.1     3962.4   80065.6

Legend:
Flags:
(D)efault SRP : F = FBA Default, C = CKD Default, B = Both, . = N/A
(R)DFA DSE    : X = Usable, . = Not Used

```

Output descriptions:

- Usable — indicates the usable capacity of all the disk groups in the SRP.
- Allocated — indicates the sum of the device allocations, snapshot allocations, and SRDF/A DSE allocations on the SRP.
- Free — lists the difference between the usable and the allocated capacity.
- Default SRP — indicates if the SRP is the default for FBA devices.
- RDFA DSE — indicates if the SRP usable for SRDF/A DSE operations.

Output without **-detail** option (no capacity metrics):

```

      STORAGE RESOURCE POOLS

Symmetrix ID   : 000197100087

-----
Name              Flg      Description
                   DR
-----
DEFAULT_SRP      FX      The Default SRP that will be used for all
                   applications
TEST_SRP         ..      SRP for testing

Legend:
Flags:
(D)efault SRP : F = FBA Default, C = CKD Default, B = Both, . = N/A
(R)DFA DSE    : X = Usable, . = Not Used

```

Output with -fba filter:

```

      STORAGE RESOURCE POOLS

Symmetrix ID   : 000197100087

                                     C A P A C I T Y
-----
Name              Flg      Usable  Allocated   Free   Subscribed
                   DR
-----

```

Name	DR	(GB)	(GB)	(GB)	(GB)
DEFAULT_SRP	FX	1847.4	515.1	1332.3	80065.6
Total		1847.4	515.1	1332.3	80065.6

Legend:
Flags:
(D)efault SRP : F = FBA Default, C = CKD Default, B = Both, . = N/A
(R)DFA DSE : X = Usable, . = Not Used

List Storage Resource Pools details

Example

To display a detailed list of all SRPs in array 084, enter:

```
symcfg -sid 084 list -srp -v
```

Sample output

```
Symmetrix ID          : 000197800084
Name                  : DEFAULT_SRP
Description           :
Default SRP          : FBA
Usable Capacity (GB) : 8833.8
Used Capacity (GB)   : 494.9
Free Capacity (GB)   : 8339.9
Subscribed Capacity (GB) : 32239.8
Subscribed Capacity (%) : 297
Reserved Capacity (%) : None
Compression State    : Enabled
Data Reduction Ratio : 2.2:1
Usable by RDFA DSE   : Yes

Disk Groups (2):
{
-----
#   Name                Flgs Speed FBA CKD   Usable
-----
   1 DISK_GROUP_001     IEN   N/A 100  0   8833.8 Internal
-----
Total                    100  0   8833.8
}

Available Service Levels (1):
{
  Diamond
}

Legend:
Flags:
Disk (L)ocation:
  I = Internal, X = External
(T)echnology:
  E = Enterprise Flash Drive, F = Fibre Channel,
  S = SATA, - = N/A
(S)tatus:
  N = Normal, D = Degraded, F = Failed

Name                  : SRP_2
Description           :
Default SRP          : None
Usable Capacity (GB) : 6437.8
Used Capacity (GB)   : 0.9
```

```
Free Capacity (GB)      : 6436.9
Subscribed Capacity (GB) : 0.9
Subscribed Capacity (%) : 0
Reserved Capacity (%)  : 15
Compression State      : Enabled
Data Reduction Ratio   : 2.1:1
Usable by RDFA DSE    : No
```

```
Disk Groups (1):
```

```
{
-----
#   Name                               Flgs Speed FBA CKD   Usable
   Name                               LTS  (rpm) (%) (%)   Capacity
-----
  2 DISK_GROUP_002                   IEN   N/A 100  0    6437.8 Internal
-----
Total                               100  0    6437.8
}
```

```
Available Service Levels (1):
```

```
{
  Diamond
}
```

```
Legend:
```

```
Flags:
```

```
Disk (L)ocation:
```

```
I = Internal, X = External
```

```
(T)echnology:
```

```
E = Enterprise Flash Drive, F = Fibre Channel,
```

```
S = SATA, - = N/A
```

```
(S)tatus:
```

```
N = Normal, D = Degraded, F = Failed
```

Output descriptions:

- Default SRP — indicates if the SRP is the default pool for FBA devices.
- Usable Capacity — sum of the useable capacity of all the disk groups in the SRP.
- Allocated Capacity — sum of the device allocations, the snapshot allocations, and the SRDF DSE allocations on the SRP.
- Free Capacity — difference between the useable and the allocated capacity.
- Subscribed Capacity — sum of the sizes of all the thin devices subscribed against the SRP.
- Subscribed Capacity percentage — percentage of the Usable Capacity used by thin devices.
- Reserved Capacity — percentage of the Useable Capacity reserved for non-snapshot activities.

i NOTE:

Existing TimeFinder snapshot sessions created on devices in the SRP may become invalid if the Free Capacity of the SRP, as a percentage of the Usable Capacity, goes below the Reserved Capacity.

- Compression State: indicates if compression is enabled or disabled for the SRP.
- Data Reduction ratio: ratio of Logical Allocated Capacity of all devices in SRP and the Physical Allocated Capacity for all devices in SRP.
- Useable by RDFA DSE — indicates if the SRP is usable for SRDF/A DSE operations.
- Disk Groups — lists the disk groups configured to the SRP. The list of disk groups is sorted by disk group number.
- Available SLOs — lists the Service Levels that are available for this SRP.

Show specific Storage Resource Pool

Description

The `symcfg show -srp` command displays a specified Storage Resource Pool (SRP).

Options

-detail

Shows detailed information about a specific SRP.

-fba Shows capacities for SRP with FBA emulation.

-ckd Shows capacities for SRP with CKD emulation. This option is available for arrays running HYPERMAX OS 5977 Q12106SR or higher.

Example

To show detailed information about the pool named SRP_1 on array 063, enter:

```
symcfg -sid 063 show -srp SRP_1
```

Sample output

```
Symmetrix ID           : 000197200063

Name                   : SRP_1
Description            : This is Default SRP2
Default SRP           : Both
Effective Used Capacity (%) : 15
Usable Capacity (GB)  : 12915.8
Allocated Capacity (GB) : 1853.9
Free Capacity (GB)    : 11061.9
Subscribed Capacity (GB) : 152147.3
Subscribed Capacity (%) : 1177
Reserved Capacity (%) : 10
Compression State     : Disabled
Data Reduction Ratio  : N/A
Usable by RDFA DSE    : Yes

Disk Groups (2):
{
-----
#   Name                Flgs Speed FBA CKD   Usable
   Name                LTS  (rpm) (%) (%)   Capacity
-----
  1 GRP_1_300_10K_3R5   IFN 10000  80  19   12040.9 Internal
  2 GRP_2_960_EFD_R1   IEN   N/A 100   0    874.9 Internal
-----
Total                               81  18   12915.8
Available Service Levels (6):
{
  Optimized
  Diamond
  Platinum
  Gold
  Silver
  Bronze
}
Legend:
Flags:
  Disk (L)ocation:
    I = Internal, X = External
  (T)echnology:
    E = Enterprise Flash Drive, F = Fibre Channel,
    S = SATA, - = N/A
  (S)tatus:
    N = Normal, D = Degraded, F = Failed
```

Storage Resource Pool demand reporting

The `symcfg list -srp -demand` command reports the subscribed and the allocated demand, from Service Levels and Storage Groups, placed on the Storage Resource Pools (SRPs).

Report Service Level demand on Storage Resource Pools

Options

-type sl

Reports the demand placed by Service Levels on the Storage Resource Pools (SRP)s. Devices that are not in a FAST managed storage group or are in a FAST managed storage group without an explicit Service Level set will be reported against the Optimized Service Level.

-details

Reports the details of the Service Level demand on the SRP.

-fba

Reports Service Level demand on the SRP for FBA emulations.

-ckd

Reports Service Level demand on the SRP for CKD emulations. This option is available for arrays running HYPERMAX OS 5977 Q12016SR or higher.

Example

To report detailed Service Level demand on the SRPs for -fba emulations, enter:

```
symcfg -sid 063 list -srp -demand -type sl -fba
```

To report detailed Service Level demand on the SRPs for -ckd emulations, enter:

```
symcfg -sid 063 list -srp -demand -type sl -ckd
```

Sample output

 **NOTE:** The SRDF DSE allocations come from the FBA default pool and will report as 0.0 when using the -ckd filter.

Output with -fba filter:

```
STORAGE RESOURCE POOLS

Symmetrix ID : 000197200063

Name           : SRP_1
Usable Capacity (GB) : 10531.6
SRDF DSE Allocated (GB) : 0.0
Snapshots Allocated (GB) : 0.0

-----
Service Level      Subscribed      Allocated
Name              (GB)  (%)      (GB)  (%)
-----
<none>            152028.3 1443      1746.1  1
Gold              112.8    1         107.8  95
Optimized         0.0      0         0.0    0
-----
Total             152141.2 1444      1853.9  1
```

Output with -ckd filter:

```
STORAGE RESOURCE POOLS

Symmetrix ID : 000197200063

Name           : SRP_1
Usable Capacity (GB) : 2384.2
SRDF DSE Allocated (GB) : 0.0
Snapshots Allocated (GB) : 0.0

-----
Service Level      Subscribed      Allocated
Name              (GB)  (%)      (GB)  (%)
-----
```

<none>	6.2	0	0.0	0
Total	6.2	0	0.0	0

Report storage group demand on Storage Resource Pool

Options

- type sg**
Reports the demand placed by storage groups on the Storage Resource Pools (SRP)s.
- fba**
Reports SG demand on the SRP for FBA emulations.
- ckd**
Reports SG demand on the SRP for CKD emulations. This option is available for arrays running HYPERMAX OS 5977 Q12016SR or higher.

Examples

To report SRP capacity, enter:

```
symcfg list -demand -srp -sid 188
```

To report detailed SRP capacity information, enter:

```
symcfg list -demand -detail -srp -sid 188
```

To report storage group demand on the SRPs, enter:

```
symcfg -sid 188 list -srp -demand -type sg
```

Sample output

```

          S R P      C A P A C I T Y      R E P O R T
Symmetrix ID   : 000197800188
-----
                                     Subscribed Capacity  Snapshot Capacity  Usable Capacity
                                     -----
Name                                     Total  Allc      Total Mdfy      Total  Used
                                     (TB)  (%)      (TB)  (%)      (TB)  (%)
-----
SRP_1                                     208.00  77      416.00  14      116.30  66

```

```

          S R P      C A P A C I T Y      R E P O R T
Symmetrix ID   : 000197800188
Usable Capacity
-----
Used      User  System  Temp      Total Allc NonShared  Shared  Total Mdfy NonShared  Shared  Total
Name      (TB)  (TB)   (TB)      (TB)  (%)   (TB)  (TB)   (TB)  (%)   (TB)  (TB)   (TB)
(%)
-----
SRP_1    74.20  2.0   22.30      208.00  77   144.00  16.00  416.00  14   41.60  14.60  116.30
66

```

```

STORAGE RESOURCE POOLS
Symmetrix ID   : 000197800188
Name           : SRP_1
Usable Capacity (GB) : 10062.0
Subscribed Capacity(%) : 67.6
Compression State : Enabled
Compression Ratio : 2.0:1

```

SRDF DSE Allocated (GB) : 0.0

SG Name	Subscribed (GB)	Allocated (GB) (%)	Used (GB)	Comp Ratio	Allocated (GB)	Snapshot Used (GB)	Snapshot Comp Ratio
SG1	60.0	50.0 83	25.0	2.0:1	1.0	1.0	1.0:1
SG2	20.0	0.0 0	0.0	-	0.0	0.0	-
SG3	6723.1	1291.0 19	645.5	2.0:1	2.0	1.6	1.3:1
<not_in_sg>	25.0	20.0 80	20.0	1.0:1	0.0	0.0	-
Total	6828.1	1361.0 20	690.5	2.0:1	3.0	2.6	1.2:1

List Storage Resource Pools for thin devices

Options

-fba

Reports thin devices with FBA emulations.

-ckd

Reports thin devices with CKD emulations. This option is available for arrays running HYPERMAX OS 5977 Q12106SR or higher.

-mb or -tb

Reports capacities in MB (megabytes) or TB (Terabytes). Default capacities are reported in GB (Gigabytes).

Examples

To list thin devices and associated SRPs, enter:

```
symcfg list -tdev -srp
```

To list thin devices with `-ckd` emulation, and associated SRPs, enter:

```
symcfg list -tdev -srp -ckd
```

To list thin devices with `-fba` emulation, and associated SRPs, enter:

```
symcfg list -tdev -srp -fba
```

Sample output

Output with no emulation type filter:

```
Symmetrix ID : 000197100087
```

S Y M M E T R I X T H I N D E V I C E S

Sym Dev	Flgs EC	Thin Device			Snapshots	
		Total (GB)	Allocated (GB)	(%)	Allocated (GB)	SRP Name
00001	FX	3.2	0.0	0	0.0	DEFAULT_SRP
00002	FX	3.2	0.0	0	0.0	DEFAULT_SRP
...						
00105	9X	0.9	0.0	0	0.0	DEFAULT_SRP
00106	9X	0.9	0.0	0	0.0	DEFAULT_SRP
00107	9X	0.9	0.0	0	0.0	DEFAULT_SRP
...						
010E0	FX	0.2	0.0	0	0.0	DEFAULT_SRP
010E1	FX	0.2	0.2	100	0.0	DEFAULT_SRP
Total		80066.7	515.5	0	0.0	

Legend:
Flags:


```

Device (E)mulation  A = AS400, F = FBA, C = CKD, 8 = CKD3380, 9 = CKD3390
(C)urrent SRP      X = Device is currently associated with SRP, . = N/A

```

Output descriptions:

- Total — configured size.
- Allocated — device allocations.
- Snapshots Allocated — snapshots allocated to SRP.
- SRP Name — associated Storage Resource Pools (SRP)s.

NOTE: A device only has allocations on more than one Storage Resource Pool, if the device is in the process of moving its allocations to a target Storage Resource Pool. (For example, the pool specified on the storage group that contained the device was changed to a different Storage Resource Pool).

Output with -ckd filter:

```

Symmetrix ID : 000197100087

          S Y M M E T R I X   T H I N   D E V I C E S
-----
Sym  Flgs      Thin Device      Snapshots
Dev  EC        Total   Allocated  Allocated
              (GB)      (GB) (%)      (GB) SRP Name
-----
00105  9X          0.9         0.0  0         0.0 DEFAULT_SRP
00106  9X          0.9         0.0  0         0.0 DEFAULT_SRP
00107  9X          0.9         0.0  0         0.0 DEFAULT_SRP
00108  9X          0.9         0.0  0         0.0 DEFAULT_SRP
00109  9X          0.9         0.0  0         0.0 DEFAULT_SRP
0010A  9X          0.9         0.0  0         0.0 DEFAULT_SRP
0010B  9X          0.9         0.0  0         0.0 DEFAULT_SRP
0010C  9X          0.9         0.0  0         0.0 DEFAULT_SRP
0010D  9X          0.9         0.0  0         0.0 DEFAULT_SRP
0010E  9X          0.9         0.0  0         0.0 DEFAULT_SRP
-----
Total                8.8          0.0  0         0.0

Legend:
Flags:
Device (E)mulation  A = AS400, F = FBA, C = CKD, 8 = CKD3380, 9 = CKD3390
(C)urrent SRP      X = Device is currently associated with SRP, . = N/A

```

Output with -fba filter:

```

Symmetrix ID : 000197100087

          S Y M M E T R I X   T H I N   D E V I C E S
-----
Sym  Flgs      Thin Device      Snapshots
Dev  EC        Total   Allocated  Allocated
              (GB)      (GB) (%)      (GB) SRP Name
-----
00001  FX          3.2         0.0  0         0.0 DEFAULT_SRP
00002  FX          3.2         0.0  0         0.0 DEFAULT_SRP
...
010E0  FX          0.2         0.0  0         0.0 DEFAULT_SRP
010E1  FX          0.2         0.2 100        0.0 DEFAULT_SRP
-----
Total                80057.9       515.5  0         0.0

Legend:
Flags:
Device (E)mulation  A = AS400, F = FBA, C = CKD, 8 = CKD3380, 9 = CKD3390
(C)urrent SRP      X = Device is currently associated with SRP, . = N/A

```

Manage Configuration Changes

This chapter describes configuration change concepts and explains how to perform array and device change operations.

Topics:

- Configuration change management overview
- Verify and check sessions
- Abort configuration session
- Configuration change guidelines
- Supported configuration operations
- Set array attributes
- Set Service Level name
- Customize Service Level Response Time Multiplier
- External disk group management
- Create devices (HYPERMAX OS 5977 Q12016SR or higher)
- Create devices (HYPERMAX OS 5977 lower than 5977 Q12016SR)
- Expand device capacity
- Copy devices
- Convert devices
- Set device attributes
- Set device identifiers
- Delete devices
- Device reservation management
- Device pool management
- Map CKD devices to CU image (HYPERMAX OS 5977 Q12016SR or higher)
- Unmap CKD devices from CU image (HYPERMAX OS 5977 Q12016SR or higher)
- Assign PAV alias addresses to CU image mapped devices
- Remove PAV alias addresses from CU image
- Map FBA devices to director ports
- Unmap FBA devices from director port
- I/O activity and unmapping devices
- Managing PowerPath initiator and host registration
- Introduce devices to a host
- Set SRDF group attributes
- Swap RA groups
- Virtual Witness (vWitness)
- Add director
- Remove director
- Port to director emulation support
- Set port characteristics

Configuration change management overview

Use the SYMCLI `symconfigure` command or the `symdev` command to make configuration changes to a locally-connected array or to an RDF-linked array. This command is run from the local host and performs control operations on arrays, as well as array devices, groups, directors, and ports.

Array controls include:

- Setting array-wide metrics
- Determining what type of devices the array will support, such as RAID 6 devices

Device controls include:

- Creating, modifying, and deleting devices
- Mapping and masking devices
- Configuring device pools
- Reserving devices
- Releasing device reservations

Configuration change session rules

The following rules apply to change sessions and concurrent change sessions:

- Concurrent Provisioning — Up to four concurrent configuration change sessions can run at the same time, when they are non-conflicting. Multiple parallel configuration change sessions can run at the same time as long as the changes do not include any conflicts on the following:
 - Device back-end port
 - Device front-end port
 - Device
- The array manages its own device locking.
- A session ID identifies each running session on the array.

Symconfigure command execution formats and options

Description

Configuration changes that are submitted to the array are processed in a session. The `symconfigure` command has several formats for executing array configuration changes. The command file format can contain various command entries terminated with a semicolon (;). Multiple changes can be specified in one session, but all changes must fall into one complete operation — for example, creating a device, adding the device to a device pool, and enabling the device state. Additional command file examples are provided with the actions described later in this chapter.

Refer to the Dell EMC Solutions Enabler CLI Reference Guide for the complete manpage description of the `symconfigure` command.

Options

abort

Gains control of an existing session to abort it and free any held locks.

commit

Attempts to apply the changes defined in the command file into the specified array.

list

Lists the relevant details, depending on the option:

- `-freespace` shows the free physical disk space within the array as it can be used to create new devices for different emulation modes. Free disk space on unformatted disks is shown as available for all emulation modes. If a physical disk has been partially used to create a device, that device is considered formatted and the rest of the available space can only be used for devices of the same emulation mode.
- `-v` displays configuration information that is not stored in the SYMAPI database and that needs to be retrieved directly from the configuration server.
- `-reserved` shows a summary of all reservations.

prepare

Validates the syntax and correctness of the operations. Verifies the validity of the command file changes and their appropriateness for the specified array. The prepare action has no function for pool sessions.

preview

Ensures the command file syntax is correct. Verifies the validity of the command file changes.

query

Returns information about the status of a configuration change session.

release

Releases the specified device reservation.

reserve

Processes the command file to reserve the indicated devices and displays the resulting reserve ID.

show

Shows the details of the specified device reservation.

verify

Verifies that the configuration currently running in the specified array complies with the requirements for host-based configuration changes.

Examples

There are various ways to execute the `symconfigure` command.

- Using the `-file` option:

```
symconfigure commit -sid 3160 -file unmap_dev.cmd
```

Where `unmap_dev.cmd` contains:

```
unmap dev 00020:00024 from dir ALL:ALL;
```

NOTE: When using the `symconfigure -file` option, text files can have a maximum comment of 512 characters on Windows. Make sure the comment line does not exceed 512 characters.

- Redirect a number of screen entries to `stdin` to save keystroke entries (for UNIX platforms), and not use a command file.

For example, to prepare a chain of `symconfigure` commands on the screen to be redirected to `stdin`, enter:

```
symconfigure -sid 1234 prepare <<DELIM
create dev count=3 size=3200 cyl,
emulation=FBA, config=2-Way-Mir;
create dev count=1, size = 3200 cyl,
emulation=FBA, config=unprotected;
DELIM
```

- Use the `-cmd` option. With this option, the commands that would normally be put in a command file are enclosed in quotes. A command can run over many lines, but you cannot press **Enter**. For example:

```
symconfigure -sid 256 -cmd "create dev count=3, size = 3200 cyl, emulation=FBA, config=2-Way-Mir;create dev count=1, size = 3200 cyl, emulation=FBA, config=unprotected;" -v -nop preview
```

Verify and check sessions

Description

Use the `symconfigure verify`, `preview`, and `prepare` arguments to verify and check a change session that is acting on a specified `symconfigure` command file. The `commit` argument performs these same checks, then attempts to execute the specified configuration change.

The `symconfigure` command modifies configuration operations and components. A lock may be taken out by the specified VMAX array configuration server while the configuration change session is active. Use the `query` option to monitor the stages of session processing. Not all stages are always executed. Use caution when controlling which stages are to be completed, to allow checking and debugging of the command file before the changes are implemented.

Syntax

To verify a change session, use the following syntax:

```
symconfigure -sid <SymmID> verify
preview
prepare
commit
```

```
symconfigure -sid <SymmID> [-i <Interval>] [-c <Count>]
[-session_id <SessionID>] [-v]query
```

Options

verify

Determines if a VMAX array can be modified, as there are restrictions as to what state the configuration must be in before allowing changes to be applied.

preview

Verifies that each individual change is valid and the syntax is correct, and then terminates the session without change execution.

prepare

Performs the preview checks, validates the change operation (devices are in correct state, etc.), then terminates the session without attempting to make the configuration change.

commit

Completes all checks and verifications, and then attempts to make the requested configuration changes in the specified array.

query

Checks the status of any configuration change session or whether there is a configuration session running. This option is useful in SRDF environments, where a change to a local array on one host results in a corresponding change to a remote array. The System Manager of a host, connected to the remote array, can monitor the progress of the change. A `query` is also helpful at sites where the Symmetrix Optimizer is modifying a configuration by rearranging the placement of data.

Examples

To check the status of change session 100, on array **345**, every 10 seconds for the next two minutes, enter:

```
symconfigure -sid 345 query -i 10 -c 12 -session_id 100
```

Abort configuration session

Options

-session_id

Aborts specified session when multiple sessions are running on the array. If `-session_id` is not specified the abort command displays each running session and prompts for the session ID.

Examples

The `abort` option allows you to stop a configuration session. To abort a change session on a specific array, enter:

```
symconfigure -sid 12345 abort
```

To abort change session 100, enter:

```
symconfigure -sid 343 abort -session_id 100
```

Restrictions

- Because changes made in the SRDF operations class will initiate actions on the local and remote arrays, it might become necessary to abort processing on a remote array.

- At some point during `commit` processing, a *point of no return* is reached. Any attempt to abort will be denied once processing has reached or gone beyond this point.

Configuration change guidelines

Understand your array configuration before making configuration changes. Any changes can impact stored data.

i **NOTE: If you have problems with a new configuration, contact the EMC Customer Support Center for assistance in reverting to your previous configuration.**

Ensure that all critical data is preserved and safe when creating new or changing device configurations, and do not store data on any device that is not mirrored, or RAID-protected. All configuration changes and device attribute adjustments must meet certain open systems guidelines detailed in the E-Lab Interoperability Navigator, which can be viewed at <http://elabnavigator.EMC.com>.

Verify viable array configuration

Syntax

Verify that the current array configuration is a viable configuration for host-initiated configuration changes. To verify the current array configuration is ready for changes, use the following syntax:

```
symconfigure verify -sid SymmID
```

Check for free physical disk space

Syntax

Before creating new devices, check for free physical disk space using the following syntax:

```
symconfigure list -freespace [-units CYL | MB] -sid SymmID
```

i **NOTE: Free disk space on unformatted disks is shown as available for all emulation modes. New devices are created first on physical disks that have no prior allocations, causing these disks to be committed to that emulation type.**

Examine the distribution of free space across formatted disks to see if the desired mirroring can be provided, using one of the following syntax formats:

```
symdev list -sid SymmID -da all -space
```

```
symdisk list -sid SymmID
```

Stop I/O activity on affected devices

Description

Configuration changes begin only after issuing the `commit` action. Some classes of change operations may or may not impact current I/O. When possible, before issuing the `commit` action, stop I/O activity on the affected devices.

i **NOTE: If I/O activity on an affected device occurs before or during a `commit` action, the action may fail. At the very least, heavy I/O activity on unaffected devices impacts how long it takes to commit changes.**

Syntax

If required, use the following syntax to set the impacted devices for change to be *Not Ready*:

```
symdmg -g DgName not_ready [LdevName [LdevName... ]]
```

Mixed array environments

Before issuing a `symconfigure` command to an array containing both open system and mainframe devices, verify that the mainframe Missing Interrupt Handler (MIH) period for each mainframe server attached to the array is set to at least two minutes.

- To view the current MIH period, use the z/OS command `D IOS,MIH`, and note the value for the DASD device class (for example, `DASD=1:30`)
- To change the MIH period, use the z/OS command `SETIOS MIH,DASD=mm:ss`, where `mm:ss` is a period of time in minutes and seconds. Then use the z/OS command `D IOS,MIH` to verify your changes.

Once you have completed your configuration change session, use the `SETIOS MIH,DASD=mm:ss` command to set the MIH period back to its original value. For more information, consult the mainframe system administrator.

Update host with device mapping information

About this task

After a configuration change, the host's device information must be updated. Attempting host activity with a device after it has been removed or altered, but before updating the host's device information, can cause host errors. To update the host:

Steps

1. Commit a `symconfigure map` operation.
2. Run the utilities specified for the host platform as described in [Introduce devices to a host](#) on page 199.
3. Issue the `symcfg discover` command to update the SYMAPI database with the new device mapping information.
4. Resume I/O activity.

Supported configuration operations

HYPERMAX OS 5977 supports the following configuration operations.

Create device	FBA Celerra® FBA CKD AS400_D910_099 Thin devices Copy device Add gatekeeper ACLX SCSI3_persist_reserv DIF1 AS400_GK Set device identifiers Delete device
Convert director type	FA to RDF/RDF to FA Add new emulation type

	Remove emulation type
Change device emulation	FBA Celerra FBA
Thin Provisioning	Create thin devices Rename thin pool View thin pools
Map device (only on non-ACLX enabled ports)	AS400 FBA CKD devices to CU image
Unmap device	AS400 FBA CKD devices from CU image
Set array-wide parameters	Set front-end port attributes RDFA cache percent RDFA host throttle time
SRDF operations	Add RDF mirror Remove RDF mirror
Device pools	Rename pool
FAST.X	Add external disk Add external disk group Remove external disk group

Set array attributes

Syntax

Use the following command to set certain attributes to control array behavior:

```
symconfigure set Symmetrix [MetricName=MetricValue, MetricName=MetricValue];
```

Options

MetricName

The metric to be set.

Table 12. Possible metric names (HYPERMAX OS 5977)

<code>rdfa_cache_percent</code>	Sets the percentage of write pending cache that can be used by SRDF/A. This is a value from 0 to 100 percent.
<code>rdfa_host_throttle_time</code>	Sets the number of seconds to throttle host writes to SRDF/A devices when cache is full, before dropping RDF/A sessions. Throttling delays a write from the host until a cache slot becomes free. Values are from 0 to 65535.

View the array metrics

Syntax

To view the current settings for array metrics, use the following syntax:

```
symcfg -sid SymmID -v list
```

Set Service Level name

Syntax

To set to a user-defined Service Level name or reset to the default name, use the following syntax with `symconfigure`:

```
set sl <SLName=NewSLName|BASE NAME>
```

Examples

To set the Service Level default name to a user-defined name, enter:

```
symconfigure -sid 234 commit -cmd "set sl Platinum name=Emc_Platinum;"
```

```
Execute a symconfigure operation for symmetrix '000000001234' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
  Establishing a configuration change
  session.....Established.
  Processing symmetrix 000000001234
  {
    set sl Platinum name=Emc_Platinum;
  }
  Performing Access checks .....Allowed.
  . . .
  Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

To reset the Service Level default name to the default name, enter:

```
symconfigure -sid 234 commit -cmd "set sl Emc_Gold name=BASE NAME;"
```

Rules and restrictions for set/reset Service Level name (HYPERMAX OS 5977 or higher)

The following security privileges are required to execute this command:

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin

The following naming rules apply:

- Names cannot exceed 32 characters.
- Alphanumeric characters only.
- Hyphens (-) and underscores (_) are allowed, but not as the leading character.

This command will fail if:

- The `s1_name` value is a NULL value, or if the specified SL name does not exist.
- The new Service Level name supplied is already in use, as either a base Service Level name or a Service Level current name.

NOTE: Names must be unique when folded to upper case.

- A reset to a base Service Level is performed and a user-defined Service Level does not exist.

Customize Service Level Response Time Multiplier

Description

Service Level Response Time (RT) multiplier enables you to setup the delay caused due to Service Levels.

NOTE: This feature requires arrays running or higher.

The delay factor could be customized to the following values:

- none (1x)
- low (2x)
- default (5x)

Syntax

To customize Service Level (RT) multiplier, use the following syntax:

```
symcfg -sid <SymmID> [-noprompt] set -sl_rt_multiplier <SLRTMultiplier>
```

where:

-sl_rt_multiplier

Specifies the SL Response Time multiplier value.

Examples

To set the Service Level RT Multiplier to 1x on array 048, enter:

```
symcfg -sid 048 set -sl_rt_multiplier NONE
```

External disk group management

Solutions Enabler includes a feature called Federated Tiered Storage (FTS). With FTS, an external LUN (eDisk) can be attached through the SAN to the array and can be used as external back-end disks for that array.

Refer to [FAST.X](#) on page 266 for eDisk configuration details.

Create devices (HYPERMAX OS 5977 Q12016SR or higher)

Description

Use the `symdev create` command to create new devices, which includes both FBA and CKD devices.

The `symconfigure create` command is also used to create devices but it does not support model types for 3390 CKD devices. See [Create devices \(HYPERMAX OS 5977 lower than 5977 Q12016SR\)](#) on page 180 for syntax for this command.

NOTE: For CKD devices, CU images must exist for TDEV creation and mapping. Solutions Enabler does not support CU image creation. CU image and split management is done through SymmWin configuration change functions, and is performed by Dell EMC customer support during installation.

Valid device configurations are:

- BCV+TDEV
- TDEV

Syntax

To create one or more devices, use the following syntax:

```
symdev -sid <SymmID> [-noprompt] [-sg <SgName>]
create -tdev -cap <#> [-capttype <cyl|mb|gb|tb>] [-N <#>]
    [-bcv] [-emulation fba|ckd3380|as400|celerra]
    [-mobility]
    [-device_name <DeviceName> [-number <n | SYMDEV>]]
create -tdev -emulation ckd3390
    <<-model <1|2|3|9|27|54>> |
    <-cap <#> [-capttype <cyl|mb|gb|tb>]>>
    [-device_name <DeviceName> [-number <n | SYMDEV>]]
create <-gk|as400_gk|-pedev> [-N <#>]
    [-device_name <DeviceName> [-number <n | SYMDEV>]]
```

Options

-sg <SgName>

Adds the newly created device into the specified storage group.

-cap <#>/ -capttype (cyl/mb/gb/tb)

Specifies the size of the device needed in number of cylinders (default), megabytes, gigabytes, or terabytes.

Table 13. Maximum device sizes

	MBs	CYLs	GBs	TBs
FBA	671108864	71582788	65536	64
CKD-3380		3993		
CKD-3390		262668 - Maximum size for arrays running HYPERMAX OS 5977 Q12016SR and Q42016SR		
		1,182,006 - Requires arrays running HYPERMAX OS 5977 Q22017SR or higher.		

-N <#>

Specifies the number of devices.

-emulation

Specifies the device emulation type. Valid emulations are :

- FBA
- Celerra_FBA
- AS/400
- CKD-3380
- CKD-3390

-mobility

Specifies devices with mobility safe ID. Supports FBA (excluding Celerra FBA and AS400 D910), Gatekeeper, and PE (Protocol Endpoint) devices.

-device_name

Specifies device name. Supported for arrays running HYPERMAX OS 5977 Q22017SR or higher.

-model <1|2|3|9|27|54>

Specifies one of the CKD-3390 emulation models:

- 1 = CKD3390-1
- 2 = CKD3390-2
- 3 = CKD3390-3
- 9 = CKD3390-9
- 27 = CKD3390-27
- 54 = CKD3390-54

Examples

To create four, CKD thin devices each with a capacity of 16000 cylinders, enter:

```
symdev create -sid 005 -tdev -cap 16000 -N 4 -emulation ckd3390 -model 1
```

```
Create operation succeeded.
```

Create devices (HYPERMAX OS 5977 lower than 5977 Q12016SR)

Description

Use the `symconfigure create dev` command or the `symdev create` command to create new devices.

Syntax

To create one or more devices, use the following `symconfigure` syntax:

```
symconfigure create dev count=<n>,  
    size = <n> [MB | GB | CYL],  
    emulation=<EmulationType>,  
    config=<DevConfig>  
    [, preallocate size = <ALL>  
    [, allocate_type = PERSISTENT]]  
    [, sg=<SgName>]  
    [, mapping to dir <director_num:port>  
    [starting] target = <scsi_target>,  
    lun=<scsi_lun>, vbus=<fibre_vbus>  
    [starting] base_address = <cuu_address>[,...]]  
    [, mapping to cu_image = <cu_image_num>,  
    split_name=<split_name>,  
    [starting] base_address=<base_address>  
    [, mvs_ssid=<n>]]  
    [, device_attr =  
    <SCSI3_PERSIST_RESERV | DIF1 |  
    AS400_GK>[,...]]  
    [, device_name=<DeviceName> [,number=<n | SYMDEV> ]];
```

To create one or more devices, use the following `symdev` syntax:

```
symdev -sid <SymmID> [-noprompt]  
    create -tdev -cap <#> [-capttype <cyl|mb|gb|tb>] [-N <#>] [-bcv]  
    create <-gk|as400_gk|-pedev> [-N <#>]
```

Options

count/#

Indicates the number of devices to create.

size/-cap/-captype

Specifies the size of the device needed in number of cylinders (default), megabytes, gigabytes, or terabytes.

Table 14. Maximum device sizes (HYPERMAX OS 5977)

FBA devices	MBs	CYLs	GBs	TBs
HYPERMAX OS 5977 Q32015SR	67108864	71582788	65536	64
HYPERMAX OS 5977 lower than Q32105SR	16777216	8947848	16384	16

emulation

Specifies the device emulation type. Valid emulations are:

- FBA
- Celerra_FBA (external gateway)
- AS/400_D910_099

config

The device configuration type. Valid types are:

- TDEV
- BCV+TDEV

preallocate size

Indicates the amount of space pre-allocated to the thin device(s) when it is bound to a pool. The only valid value is ALL.

allocate_type

Indicates the allocation type. The only valid value is PERSISTENT.

sg

The local array storage group name for devices involved in RDF operations.

mapping to dir

Specifies the director and port for the mapping operation

For arrays running HYPERMAX OS 5977, `mapping to dir` returns a "feature not supported" error when the request includes mapping of devices to be created to ACLX enabled front end ports. In addition, the valid range for port is extended to 0 to 31.

scsi_target

Indicates a hex value for the SCSI target ID.

scsi_lun

Indicates a hex value for the SCSI logical unit number.

vbus

The virtual fibre bus address used when mapping to a Fibre Channel director port using volume set addressing.

base_address

Indicates a base or alias address for a device being mapped to an EA or EF port. These are mainframe ports which expect devices to be mapped in groups to form CU images.

mapping to CU image

Specifies the CU image for the mapping operation

device_attr

Specifies the attributes to be set on the new device. Possible values include:

- `SCSI3_persist_reserv` (persistent group reservation) You can set or clear the SCSI-3 persistence attribute only if the device is unmapped.
- `DIF1`
- `AS400_GK`

device_name

Specifies the user supplied name with a maximum of 64 characters including the suffix. Any character may be used for the device name except quotes (" "), which denote the start and end of input. The device name plus an

optional suffix can have a maximum of 64 characters. If using a numerical suffix, the device name will be limited to 50 characters (prefix) and the trailing numerical suffix number will be limited to 14 characters. If not using a numerical suffix, all 64 characters can be specified for the device name. The maximum starting suffix is 1000000. If setting a device name when creating RDF pairs, both RDF devices will be set to the same device name if using the SYMDEV option. Solutions Enabler does not check for uniqueness of device names.

number

Represents the user supplied number for the starting suffix. Specifying SYMDEV means that the corresponding device number will be used as the suffix.

Examples

```
symdev create -sid 005 -cap 10000 -nop -tdev -v
```

```
Create operation succeeded for devices: 01ff0.
```

Restrictions when creating and managing devices

The following rules and restrictions apply when creating or modifying devices:

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin.
- Allows for the creation of externally-visible (TDEV) devices (for FBA up to 64TB), without requiring meta devices. The creation and control of meta devices is no longer supported. Array-wide meta settings, previously used to automatically create meta devices are not supported.

NOTE: Requests to create FBA meta and CKD RAID-10 (ckd_meta) devices are blocked on these arrays.

- Creation of standard (disk group) provisioned devices is blocked. This includes internal devices that are used for the backing of thin pools (DATADEVs and SAVEDEVs) as well as Diskless and VDEV devices, which are implemented as disk group provisioned devices.
- Creation of gatekeeper devices is supported, but they will always be thin devices.
- The ACLX attribute is supported only for thin devices with FBA emulation.
- Encapsulation of user data as disk group provisioned devices is not supported.
- Partial allocation of thin devices is not supported.
- SRDF device types cannot be created or converted when:
 - Domino mode is enabled on any current SRDF pairs.
 - There are any invalid tracks on any of the current SRDF devices.
 - Concurrent SRDF is enabled on a device.
- The only emulation type supported for iSeries devices is D910_099, as it is the only iSeries emulation type that supports thin provisioned devices. The size range is a minimum of 1562 cylinders and a maximum of 1118481 cylinders.
- AS/400_D910_099 thin devices cannot be CKD metadevices, SAVE devices, or DATA devices. In addition, IBM i thin devices cannot have the following attributes:
 - RCVRPNT_TAG
 - ACLX
 - DIF1

To create gatekeeper devices, use the `symconfigure create gatekeeper` command. For HYPERMAX OS 5977 and higher, gatekeeper devices can be added to storage groups using the syntax `sg=SgName`. All information about gatekeepers can be found in the *Dell EMC Solutions Enabler Installation and Configuration Guide* and in the EMC Knowledgebase solution EMC255976 available on EMC Online Support.

NOTE: Only thin gatekeeper devices can be created on arrays running HYPERMAX OS 5977.

Expand device capacity

Description

The `symdev modify` command expands a device capacity. Refer to [Create devices \(HYPERMAX OS 5977 Q12016SR or higher\)](#) on page 178 or [Create devices \(HYPERMAX OS 5977 lower than 5977 Q12016SR\)](#) on page 180 for maximum FBA and CKD device sizes.

Device expansion of CKD 3390 devices is supported with VMAX arrays running HYPERMAX OS 5977 Q22017SR or higher. Automatic VTOC index rebuild must be enabled or be prepared to submit an ICKDSF batch job to rebuild the VTOC before the newly added space on the volume can be used.

CKD device expansion is not allowed for the following devices and configurations:

- CKD 3380
- TDATs
- CDK device marked as Soft Fenced
- Device target size equal to less than the current size
- CKD devices with SnapVX sessions

Syntax

To modify device capacity, use the following syntax:

```
symdev -sid <SymmID> [-noprompt]
      modify -tdev -cap <#> [-capttype <cyl|mb|gb|tb>] [-rdfg <RdfGrpNum>]
          -devs <<SymDevStart>:<SymDevEnd> | <SymDevName>
              [,<<SymDevStart>:<SymDevEnd> | <SymDevName>>... ]>
```

Options

- cap** Specifies the new expanded size for the device.
- rdfg** Lists devices that belong to the specified SRDF group. When used with `modify` it specifies the SRDF group associated with the SRDF devices and indicates that both sides of the SRDF pair, which are associated with the SRDF group, should be expanded.

Examples

To expand device `1fe0` on array `005` to 4TB, enter:

```
symdev modify 1fe0 -sid 005 -cap 4000 -capttype gb -nop -tdev
Modify operation succeeded for devices: 1fe0
```

To expand 4 devices on SRDF group `33`:

```
symdev modify -sid 85 -tdev -cap 1000 -capttype mb -dev 007D2:007D5 -v -rdfg 33 -nop
Symmetrix: 000197100085
Requested device(s): 007D2:007D5
Symmetrix: 000196801476
Requested device(s): 00A8A:00A8B 00AA2:00AA3

STARTING a TDEV Expand Device operation on Symm 000196801476.

Expanding devices [00A8A:00A8B] .....Done.
Expanding devices [00AA2:00AA3] .....Done.

The TDEV Expand Device operation SUCCESSFULLY COMPLETED on Symm 000196801476: 4 device(s)
```

expanded.

STARTING a TDEV Expand Device operation on Symm 000197100085.

Expanding devices [007D2:007D5]Done.

The TDEV Expand Device operation SUCCESSFULLY COMPLETED on Symm 000197100085: 4 device(s) expanded.

Modify devices operation succeeded.

Copy devices

Copying the attributes of an existing device into available disk space configures new devices based on the copied attributes. To configure new devices this way, either specify the quantity of disk space to configure into the new devices, or the number of devices to be created. The quantity of disk space represents the new space that will be available for the host to use, and not the space allocated in the array to manage the request according to the device protection requirements. In addition, copied attributes can be changed. For example, copied attributes of a standard device, can be changed to BCVs on the new device.

Copy a similar device

Syntax

To configure a device by copying a similar device, use the following syntax:

```
symconfigure configure [n.nn [MB | GB] | nn
devices] copying dev SymDevName[mapping to dir DirectorNum:PortNum
[masking hba [awwn=awwn | wwn=wwn |
iscsi=iscsi | aiscsi=aiscsi]
host_lun=lun |dynamic_lun]] ]
[,device_name=DeviceName [,number=n | SYMDEV]]
[overriding
[size=<n> [MB | GB | CYL]]
[emulation=<EmulationType>]
[config=<DevConfig>]
[data_member_count=<n>]
[mvs_ssid=<n>]
[disk_group=<n>| name:<DskGrpName>]]];
```

NOTE:

For arrays running HYPERMAX OS 5977, the configure operation is blocked if the source device specified by *SymDevName* is a *TDAT* device.

n.nn [MB | GB]

Specifies the quantity of disk space to configure.

nn devices

Specifies the number of devices to create.

SymDevName

The device name of the model device.

DirectorNum PortNum

The mapping attributes of the device are not copied. If the new devices are to be mapped, you must specify the director/port addresses.

mapping to dir

Specifies the director and port for the mapping operation.

Returns a feature not supported error when the request includes mapping of devices to be created to ACLX enabled front end ports. In addition, the valid range for *PortNum* is extended to 0 to 31.

masking hba

The masking attributes of the model device are not copied. If the new devices are to be masked, you must specify the host HBA to which the devices should be masked.

host_lun

Specifies the LUN addresses to be used for each device that is to be added for the host HBA.

dynamic_lun

Specifies to use the dynamic LUN addressing features but does not require a LUN address for each device. The LUN addresses are assigned based on what may already be in use for that host HBA.

device_name

Specifies the user supplied name with a maximum of 64 characters including the suffix. The legal characters for the device name include all alpha, numeric, underscore (_) and period (.). The device name plus an optional suffix can have a maximum of 64 characters. If using a numerical suffix, the device name will be limited to 50 characters (prefix) and the trailing numerical suffix number will be limited to 14 characters. If not using a numerical suffix, all 64 characters can be specified for the device name. The maximum starting suffix is 1000000.

number

Represents the user supplied number for the starting suffix. Specifying SYMDEV will mean that the corresponding device number will be used as the suffix.

overriding

Indicates that you will be overriding some of the characteristics of the copied device.

size

Specifies the size of the new devices.

emulation

Specifies the device emulation type. Valid emulations are FBA and AS/400_D910_099.

config

The device configuration type. Valid configuration types are TDEV and TDEV+BCV.

mvs_ssid

When creating devices in an array that also contains CKD devices, a z/OS (MVS) subsystem ID (**mvs_ssid**) value can be provided so the new FBA devices are not seen as part of an existing subsystem ID group.

name

Specifies the name of the remote disk group. By default, the disk group name is DISK_GROUP_xxx, where xxx, is the disk group number.

Usage is: disk_group=# or disk_group=name:DskGrpName

Convert devices

Syntax

To convert a device configuration type, use the following syntax:

```
symconfigure convert dev <SymDevName>[:<SymDevName>] to
<DevConfig> [emulation=EmulationType,]
[ ra_group=<n>, remote_dev=<SymDevName>,
invalidate=<invalidate_opt>, [remote_mvs_ssid=<n>],
start_copy=<YES | NO> ] [mvs_ssid=<n>] [raidset = [TRUE | FALSE]];
```

NOTE: An IBM i thin device (TDEV) can be converted to a BCV+TDEV device, and a BCV+TDEV device can be converted to a TDEV device.

Options

SymDevName

Specifies the name of the device targeted for change. To target more than one device, indicate the first and last devices in a series separated by a colon (:).

DeviceConfig

Specifies the desired device configuration type. [Table 15. Thin device conversions](#) on page 187 lists allowable configuration conversions.

emulation

Indicates the device's emulation type.

ra_group

Specifies the RA group number in the SRDF environment.

remote_dev

Specifies the name of the remote array device targeted for change. If a range of *SymDevNames* is specified in the first line of the convert statement, the remote *SymDevName* value is increased incrementally to arrive at the corresponding device number.

invalidate_opt

Indicates the SRDF device to invalidate so a full copy can be initiated from the remote mirror. Allowed values are R1 (invalidate the source), or R2 (invalidate the target). The value NONE is not supported in Solutions Enabler V7.0 and higher.

remote_mvs_ssid

Specifies the remote z/OS (MVS) subsystem ID that is assigned to any device created as a result of removing any mirror(s). If not provided, the original MVS SSID is assigned when available. If the MVS SSID group is full, supply a new MVS SSID. Only one `remote_mvs_ssid` can be used in a session; it is applied to all devices created within that session.

start_copy

Indicates whether an SRDF pair should be synchronized after the configuration change is committed.

NOTE:

When creating SRDF devices, all conversions within a session must have:

- **Device configuration settings that reflect the same destination SRDF type (RDF1 or RDF2).**
- **The same `ra_group` number.**
- **The same `invalidate` option.**
- **The same `start_copy` option.**

mvs_ssid

Specifies the z/OS (MVS) subsystem ID that is assigned to any device created as a result of removing any mirror(s). If not provided, the original MVS SSID is assigned when available. If the MVS SSID group is full, supply a new MVS SSID. Only one `mvs_ssid` can be used in a session; it is applied to all devices created within that session.

raidset

When requesting to convert a RAID-S group to unprotected devices, set `raidset` equal to `TRUE` and list the first RAID-S member. It is not necessary to list the other members.

Example

To convert two existing BCV devices (0001C and 0001D) to an RDF1-BCV configuration and to invalidate the source R1 and synchronize the SRDF pair, enter:


```
symconfigure convert dev 0001C:0001D to RDF1-BCV, ra group=1, remote_dev=0001c,  
invalidate=R1, start_copy=YES;
```

Device conversion restrictions

Restrictions

The following restrictions apply when converting devices:

- The `convert` command can be used for three different classes of device configuration changes, as long as the class types are performed in separate sessions.

 **NOTE: Changes for multiple operation classes can be executed in the same session, except for dynamic SRDF changes and device pool changes.**

The three class types that cannot be used in the same session are:

- Add/remove BCV/DRV attributes
- Add/remove SRDF attributes
- Increase/decrease mirroring
- Full swap operations require the R1 and R2 devices to be the same size.
- One member of a raidset cannot be converted to unprotected without converting all 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.
- The SRDF mode on an SRDF device pairing will be Adaptive Copy Disk by default, unless the SYMAPI_DEFAULT_RDF_MODE is set in the option file. If the device being converted is a diskless R1 device, the RDF mode will default to *Adaptive Copy Write Pending*, regardless of the option file setting.

Valid thin device conversions

Thin devices can be converted to other thin device configurations, as shown in the following table.


 **NOTE:**

Table 15. Thin device conversions

Original thin device	Converted to
TDEV ^a	RDF1+TDEV
TDEV ^a	RDF2+TDEV
TDEV	BCV+TDEV
BCV+TDEV ^b	TDEV
RDF1+TDEV ^a	TDEV
RDF2+TDEV ^a	TDEV
BCV+TDEV ^a	RDF1-BCV+TDEV
BCV+TDEV ^a	RDF2-BCV+TDEV
RDF1+TDEV ^a	RDF1-BCV+TDEV
RDF2+TDEV ^a	RDF2-BCV+TDEV
R1-BCV+TDEV ^a	R1+TDEV
R2-BCV+TDEV ^a	R2+TDEV
R1+TDEV	R1-BCV+TDEV
R2+TDEV	R2-BCV+TDEV

a. Not for VMAX 10K Series systems; all devices are dynamic RDF capable by default.

b. BCV+TDEV to TDEV is the only supported conversion for a BCV+TDEV device.

Convert a thin device

Example

To convert a thin device TDEV to an RDF1 thin device, enter:

```
symconfigure convert dev 00015 to RDF1+TDEV;
```

To create a thin R1 BCV device by converting a thin BCV device, enter:

```
symconfigure -sid 397 -nop -v -cmd "convert dev 015F2 to RDF1-BCV+TDEV ra_group=10
remote_dev=16A invalidate=R1 start_copy=no;" commit
```

Thin device recommendations to maximize I/O activity

Adhere the following restrictions/conditions to avoid impact on I/O activity:

- The BCV attribute is not allowed on thin SRDF devices or thin dynamic SRDF devices.
- The device being converted must not be part of a clone session.

Set device attributes

Syntax - symconfigure

To set the device attributes or emulation of a number of devices in a range, use the following form of the `symconfigure set` command:

```
symconfigure set dev SymDevName[:SymDevName]
                [emulation=EmulationType]
                [identity = NO identity]
                [attribute=[NO] device_attr];
```

Options - symconfigure

emulation

Specifies the device emulation type, which can be the following: **FBA** or **CELERRA_FBA**.

 **NOTE: You cannot set the device emulation type for Data Domain devices.**

identity

Restores the devices identity to its original value.

attribute

Indicates if a device attribute restricts how a device can be accessed.

Examples - symconfigure

To convert five devices (00015 to 00019) to Celerra FBA emulation, enter:

```
symconfigure set dev 00015:00019 emulation=CELERRA_FBA;
```

Syntax - symdev

To set the device attributes or emulation of a number of devices in a range, use the following form of the `symdev set` command:

```
symdev -sid SymmID set SymDevName <-as400_gk | -bcv | -dif1>
      -device_name DeviceName [-number <n | SYMDEV>]
      -emulation < fba | celerra >
```

To unset a device name for a device, use the `symdev unset` command:

```
symdev -sid SymmID unset SymDevName <-as400_gk | -bcv | -dif1>
```

Options - symdev

device_name

Specifies the *DeviceName* to be set.

emulation

Specifies the device emulation type, which can be the following: **FBA** or **CELERRA**.

Device attribute values

Possible values include:

- **SCSI3_persist_reserv** (From PowerMaxOS 5978, this is enabled by default for thin devices and it is not allowed to be modified)
- **AS400_GK**
- **DIF1**
- **BCV**

NOTE:

You cannot set device attributes for Data Domain devices.

Set device attribute restrictions

The following restrictions apply when setting device attributes:

- A device that is mapped or masked to an FCoE port cannot have the RCVRPVT_TAG attribute.
- When setting the device emulation type, the devices must be unmapped. No I/O to the devices involved.
- When setting the attribute type to a mapped device, it is recommended that you minimize the I/O activity to the affected devices.

The following restrictions apply to setting the DIF1 device attribute:

- The DIF1 attribute can only be set on FBA devices.
- DIF1 attribute can be set on both standard and thin host-accessible devices. You cannot set the DIF1 attribute on any internal devices.
- A device must be unmapped if resetting the DIF1 attribute.
- A device with the DIF1 attribute can only be mapped to fiber front-end directors (no iSCSI or FCoE).
- Metadevices with the DIF1 attribute must have the same state, either all set or all reset, on the metahead and all metamembers.
- DIF1 attribute can not be set on DATA and SAVE devices.
- Devices can have either the RDB_Checksum attribute or the DIF1 attribute, not both. The DIF1 flag cannot be set on a device with an active double checksum.
- Devices can have either ACLX attribute or the DIF1 attribute, not both.
- There is no relation between the DIF1 attribute and replication. Both source and target devices of any replication can have their own DIF1 setting.
- The DIF1 attribute needs to be set before requesting a reset. If the reset request is for a device range, and any one of the devices does not have the DIF1 attribute set, an error returns.
- Device emulation of a CELERRA FBA device cannot be changed to FBA if the device has AS400 GK attribute set.
- AS400_GK attribute cannot be modified if device is mapped to a port.

Set device identifiers

Description

Solutions Enabler supports device identifier management on arrays. This support allows defining names and identifiers for EMC array devices, HP devices, and VMS devices. The definitions are configured in a command file and processed with the `symconfigure` command.

A device name or identifier can be set on a single device, a range of devices, or multiple ranges of devices. Device identifiers do not have to be unique for all devices.

Syntax

The format of the command file for setting device identifiers follows:

```
symconfigure set dev SymDevName[:SymDevName]  
  [[device_name = 'DevName'] | [NO DevName]]  
  [[hp_identifier = 'hp_id'] | [NO hp_identifier]]  
  [[vms_identifier = 'vms_id'] | [NO vms_identifier]];
```

Restrictions

The device identifier can include a `device_name` and the `hp_id` or a device name and the `vms_id`. The device identifier cannot include both the `hp_id` and the `vms_id`.

- Restrictions for `device_name`:
 - `DevName` can be less than or equal to 64 characters in length.
 - Any character may be used except quotes, which are used to mark the start and end of the input.
 - The names are case sensitive.
 - There is no default device name for a device.
- Restrictions for `hp_id`:
 - `hp_id` can be less than or equal to 128 characters in length.
 - Any character may be used except quotes which are used to mark the start and end of the input.
 - The identifiers are case sensitive.
 - There is no default HP identifier for a device.
- Restrictions for `vms_id`:
 - A `vms_id` can be a number between 0 and 32766.

Identifier exclusions

Device identifiers cannot be set on the following devices:

- Metamembers. The device identifier of the metahead device will display, if applicable.
- The following internal devices (device identifiers display as N/A):
 - VAULT
 - SFS
 - DRV
 - SAVE (device names will be allowed for this device type)
 - DATA (device names will be allowed for this device type)

View device identifiers

Use the `symdev list` command to view device identifiers.

CLI commands are usually limited to an 80-character output, but if the user requests the `hp_id` of a device to be displayed, the line could be over 80 characters in length. A new form of the `symdev list` command displays the device identifiers.

This CLI command does not work in offline mode.

NOTE: Device name, device nice name, HP device identifiers, and VMS device identifiers cannot be used in any control command. All these device identifiers are only for display. Both HP device identifiers and VMS device identifiers can be set on any host and any devices.

Delete devices

Description

Use the `symconfigure delete dev` command or the `symdev delete` command to delete devices.

Syntax

To delete one or more devices, use the following syntax:

```
symconfigure delete dev SymDevName[:SymDevName]
```

```
symdev -sid <SymmID> [-noprompt]
      delete -devs <<SymDevStart>:<SymDevEnd> | <SymDevName>
              [,<<SymDevStart>:<SymDevEnd> | <SymDevName>>... ]>
```

Examples

To delete device 00015 from array 345, create a command file containing the following:

```
delete dev 00015;
```

Then commit the option using the `symconfigure` command:

```
symconfigure -sid 345 -file delete_dev.cmd -v -noprompt commit
```

Using the `symdev` command:

```
symdev delete 1f00 -sid 005 -nop
Delete devices operation succeeded.
```

Restrictions

- The `delete` commands are blocked if the source device specified by *SymDevName* is a TDAT device.
- The device must not be mapped to a front-end port.

Device reservation management

Use the configuration change functionality to reserve devices and front-end mapping addresses for future configuration and masking operations. This feature reserves the devices/addresses you plan on using, verifies that no one else has reserved the resources, and releases the reservations when the task is complete.

All reservations are assigned a reserve ID, indicating that the specified devices/addresses are reserved. Any attempt to use the reserved devices/addresses will return a message indicating that the devices/addresses are reserved.

There are two types of device reservations:

- *Enforced* — Reservations are enforced by the SYMAPI library, and require specifying the reserve ID to use the devices. This is the default behavior when reserving devices.
- *Advisory* — Reservations are enforced by co-operating applications. Some applications can ignore advisory reservations, allowing knowledgeable users to make configuration changes on reserved devices, provided that their changes are compatible with the reserving task's goal.

Both types of reservations can have expiration dates associated with them, which will automatically release a reservation if not released explicitly by the user.

Device reservations are honored only when devices are explicitly specified during Solutions Enabler configuration change operations. Operations that allow the array operating system to choose devices (such as when a meta is formed and only the metahead is specified) do not honor device reservations.

Device reservations are enabled (`TRUE`) by default in the `options` file. To disable device reservations, set the `SYMAPI_ENABLE_DEVICE_RESERVATIONS` parameter in the `options` file to `FALSE`.

Device reservations are enforced (`TRUE`) by default in the `options` file. To disable the enforcing of device reservations, set the `SYMAPI_ENFORCE_DEVICE_RESERVATIONS` parameter in the `options` file to `FALSE`.

 **NOTE:**

For information on changing the `options` file parameters, refer to the Dell EMC Solutions Enabler Installation and Configuration Guide or the Dell EMC Solutions Enabler CLI Reference Guide.

Reserve devices

Syntax

To reserve devices, use the following form:

```
symconfigure -sid SymmID [-expire ExpirationDate]  
[-f[file] CmdFile | 'redirect stdin' | -cmd "Cmd"]  
-owner Owner -comment UserComment[-enforce | -advise] reserve
```

NOTE: When reserving devices, note the returned reserve ID for subsequent processing.

Options

ExpirationDate

Specifies the date and time for a device reservation to expire.

This is an optional parameter, and if not specified, defaults to no expiration. The format for this parameter is:

`[mm/dd[/yy]:][hh:mm[:ss]]`

If you only provide the *hh:mm*, the current day will be assumed. If you only provide the *mm/dd*, the current year will be assumed. You can also specify a four-digit year.

CmdFile

Specifies the name of any ASCII text file containing a set of commands to process at a higher time.

This file can be used in the following ways:

- To reserve devices for specific configuration change operations, and the file lists configuration change commands.
- To reserve devices for non-specific operations, use the following file syntax:

```
reserve dev SymDevName[:SymDevName]
```

Using this method allows you to reserve devices for other applications.

NOTE: For arrays running HYPERMAX OS 5977, the `reserve dev` operation is blocked if the source device specified by *SymDevName* is a TDAT device.

Owner

Specifies the name of the owner of the reservation (up to 31 characters long).

UserComment

Indicates a user-specified comment detailing the device reservation (up to 255 characters long).

`-enforce`

Specifies an enforced reservation (default).

`-advise`

Specifies an advisory reservation.

Commit changes on reserved devices

Description

When committing changes on devices reserved with the `Enforced` flag, you must supply the appropriate reserve ID. If you do not have the reserve ID and someone else has reserved the devices, the commit will fail. If you have reserved the devices, or no one else has reserved the devices, then the commit will succeed.

When committing changes on devices reserved with the `Advisory` flag, some applications may not require a reserve ID. However, the `symconfigure` command does require a reserve ID.

Syntax

To commit changes on devices reserved with the `-enforce` option, use the following form:

```
symconfigure -sid SymmID -f[ile] CmdFile commit  
[-reserve_id=ResvID[,ResvID[,ResvID]]]  
[-remote_reserve_id=ResvID[,ResvID[,ResvID]]]
```

Examples

To commit the changes in the command file `delete.cmd`, enter:

```
symconfigure -sid 3241 -file delete.cmd commit -reserve_id 5
```

View reserved devices on an array

Syntax

To view the devices reserved on an array, use the following command:

```
symconfigure -sid SymmID -reserved list
```

View details on reservation ID

Syntax

To view details on a specific reservation ID, use the following command:

```
symconfigure -reserve_id ResvID show
```

Options

ResvID

The device reservation ID.

Release reserved devices

Description

When releasing device reservations, supply the appropriate reserve ID. Performing a configuration change on reserved devices will not release them. Releasing a reservation is an independent step.

Syntax

To release reserved devices, use the following form:

```
symconfigure -sid SymmID [-noprompt]  
-reserve_id ResvID[,ResvID[,ResvID]] release
```

Options

ResvID

Specifies the device reservation ID.

Examples

To release the set of devices with the ID 5 and 7, enter:

```
symconfigure -sid 3241 release -reserve_id 5,7
```

Device pool management

VMAX3 and All Flash arrays running are shipped with pre-configured thin pools. During the pre-configuration process, the physical disks are partitioned into a set of internal disk groups based on the technology, capacity, and rotational speed of the disk and protection scheme desired. In addition each disk group is fully provisioned with DATA devices, thin pools are created, and the data devices are assigned to the thin pool created for that disk group.

If physical disks are added after system delivery, the install process either creates one or more new disk groups, DATA devices, and thin pools to accommodate those disks, or adds the physical disks to existing disk pools, provisions them into DATA devices consistent with the disk group definition and adds the devices to the pool for that disk group.

For details on thin pool provisioning and restrictions, refer to [Thin device management and reporting overview](#) on page 71.

Map CKD devices to CU image (HYPERMAX OS 5977 Q12016SR or higher)

Syntax

To map a CKD device to a CU image, use the following syntax:

```
symconfigure map dev <SymDevName>[:<SymDevName>]
to cu_image = <cu_image_num>,
split_name=<split_name>,
[mvs_ssid=<n>],
starting base_address=<base_address>;
```

NOTE: Parameters after the `map dev` command can be in any order.

Options

cu_image

Specifies CU image number for mapping CKD devices to CU image.

split_name

Specifies the array subsystem split name.

mvs_ssid

MVS host subsystem ID. This option is only required if creating a new image with the mapping request.

base_address

Indicates a base address for a device being mapped to a CU image.

NOTE: The `symconfigure create dev` and `symconfigure configure dev` commands can also be used to map a CKD device to a CU image. Refer to *Dell EMC Solutions Enabler CLI Reference Guide* for command detail.

Rules and restrictions

- Requires SDR Access Type.
- Requires Storage Admin rights.
- FBA devices can only be mapped to CU images containing only FBA devices. Mixed FBA and CKD devices in a single CU image is not supported.

Unmap CKD devices from CU image (HYPERMAX OS 5977 Q12016SR or higher)

Syntax

To unmap a CKD device from a CU image, use the following syntax:

```
symconfigure unmap dev <SymDevName>[:<SymDevName>]
              from cu_image = <cu_image_num>,
              split_name=<split_name>;
```

Options

cu_image

Specifies CU image number for mapping CKD devices to CU image.

split_name

Specifies the array subsystem split name.

Rules and restrictions

- Requires SDR Access Type.
- Requires Storage Admin rights.

Assign PAV alias addresses to CU image mapped devices

Description

When assigning PAV alias addresses to CU image mapped devices, the aliases are propagated to all director ports to which the devices are mapped. Devices within the range that are not mapped are skipped. If any devices in the range are mapped to a different CU image than the first device, an error will be returned. If the device range has base addresses with gaps, the aliases will also have gaps.

Mainframe ports expect devices to be mapped in groups to form CU images. The first digit in the address is the CU image number, which can range from 0 to 0xF. The remaining two digits can range from 00 to 0xFF.

Syntax

To assign a PAV alias address range to a CU image, use the following form:

NOTE: The `add pav alias_range` command is supported on arrays running HYPERMAX OS 5977 Q12016SR or higher. It is not supported on arrays running HYPERMAX OS 5977 lower than Q12016SR.

```
symconfigure add pav alias_range nnnnn:nnnnn
              to mvs_ssid = nnn
```

To assign PAV alias addresses to devices mapped to EA or EF ports, use the following syntax:

NOTE: The `add pav alias` command is supported on arrays running HYPERMAX OS 5977 lower than Q12016SR. It is not supported on arrays running HYPERMAX OS 5977 Q12016SR or higher.

```
symconfigure add pav alias to dev SymDevName[:SymDevName], alias count=nnn
```

Examples

To add a PAV alias range to a CU image using SSID 140, enter:

```
symconfigure add pav alias_range addr 00080:0009f to mvs_ssid=140
```

To add the alias A60 to device 01D, which is already mapped, enter:

```
symconfigure add pav alias to dev 01D starting alias=A60
```

To add total PAV aliases to device 37a, enter:

```
symconfigure add pav alias to dev 37a , alias count=32
```

Remove PAV alias addresses from CU image

Syntax

To remove a PAV alias address range from a CU image, use the following syntax:

```
symconfigure remove pav alias_range from mvs_ssid =nnn
```

NOTE: The `remove pav alias_range` command is supported on arrays running HYPERMAX OS 5977 Q12016SR or higher. It is not supported on arrays running HYPERMAX OS 5977 lower than Q12016SR. The `remove pav alias` command is supported on arrays running HYPERMAX OS 5977 lower than Q12016SR. It is not supported on arrays running HYPERMAX OS 5977 Q12016SR or higher.

Map FBA devices to director ports

Description

The Device Reallocation feature allows for mapping FBA devices to front-end director ports, or mapping a range of devices to consecutive addresses starting from a specified address.

Syntax

To map a FBA device to a director port, use the following command:

```
symconfigure map dev SymDevName[:SymDevName] to dir DirectorNum:PortNum[, emulation = EmulationType]  
[starting][target = ScsiTarget,] lun = ScsiLun[, vbus = FibreVbus
```

NOTE: Parameters after the `map dev` command can be in any order.

Options

emulation

Indicates the device's emulation type. This option is required when performing operations on a Celerra device, and indicates that you are aware that you are changing the Celerra environment. Solutions Enabler supports mapping IBM i thin devices.

lun

Specifies the SCSI logical unit number (hex value).

starting

Specifies the starting address for the range of devices.

target

Specifies the SCSI target ID (hex value).

vbus

Specifies the virtual bus (vbus) address for mapping to an FA port if using volume set addressing.

Examples

To map device **00030** to director **16A**, port **0**, and SCSI target/LUN **0, 7**, enter in the command file:

```
symconfigure map dev 00030 to dir 16A:0 target=0, lun=7;
```

To map a device **00032** to director **16A**, port **0**, and SCSI target/LUN **0,2** and update the device masking database by specifying the WWN **20000000c920b484** of the host bus adapter (HBA) port through which a host accesses the device, enter the following command:

```
symconfigure map dev 00032 to dir 16A:0 target=0, lun=2, wwn=20000000c920b484;
```

Restrictions

- When mapping, there are no restrictions on I/O if adding a second path.
- After committing a `symconfigure` mapping operation, update the device mapping information within the host system environment. Attempting host activity with a device after it has been removed or altered, but before the host's device information has been updated, can cause host errors.
- To update the hosts, run the utilities designed for the specific platform as described in [Introduce devices to a host](#) on page 199. After the host environment is updated, I/O activity can resume with the array device.
- The `map dev` command returns a feature not supported error when mapping devices to ACLX enabled front end ports.
- The `map dev` command returns a feature not supported error if the devices are encapsulated devices being used as TimeFinder SnapVX source devices.

Obtain list of addresses

Syntax

To obtain a list of used addresses, including the next available address, use the following command:

```
symcfg list -SA all -address -available
```

Unmap FBA devices from director port

Description

The Device Reallocation feature allows for unmapping devices from front-end director ports.

Unmapping can be from one or all ports. Since all devices with the same SSID must either be mapped or unmapped, specify an SSID when unmapping only some devices in a CU image.

Syntax

To unmap devices from a director port, use the following form:

```
symconfigure unmap dev SymDevName[:SymDevName] from dir  
<ALL:ALL | ALL:PortNum | DirectorNum:ALL | DirectorNum:PortNum>
```

```
[, emulation = EmulationType]  
[, devmask_access = remove | retain];
```

Options

emulation

Indicates the device's emulation type. This option is required when performing operations on a Celerra device, and indicates that you are aware that you are changing the Celerra environment.

Restrictions

The `unmap dev` command returns a feature not supported error when unmapping devices from ACLX enabled front end ports.

I/O activity and unmapping devices

The following items describe how to avoid impacting I/O.

- When unmapping, no I/O activity is allowed on any devices in the specified mapped path. Devices must be made `Not Ready` or `Write Disabled`.

For example, to make the device `Not Ready`:

```
symdg create -type [REGULAR | RDF1 | RDF2] DgName symdg -g DgName -sid SymmID add dev  
SymDevName symdg -g DgName not_ready
```

- When unmapping only one path to a multi-pathed device, you may prefer to write disable that path only:

```
symdg -g DgName -SA 16A -p 0 write_disable
```

To make a device `Not Ready` without creating a device group:

```
symdev -sid SymmID not_ready SymDevName
```

NOTE: Do not use the `write_disable` argument with the `symrdf` command, as this write disables the source (R1) device(s) or the target (R2) device(s) to its/their local hosts.

- After committing a `symconfigure` mapping operation, update the device mapping information within the host system environment. Attempting host activity with a device after it has been removed or altered, but before the host's device information has been updated, can cause host errors.

To update the hosts, run the utilities designed for the specific platform as described in [Introduce devices to a host](#) on page 199. After the host environment is updated, I/O activity can resume with the array device.

Managing PowerPath initiator and host registration

The following items describe how to manage PowerPath initiator and host registrations.

- To enable PowerPath initiator registration on an array use:

```
symppath -sid <SymmID> enable -initiator_registration
```

- To disable PowerPath initiator registration on an array use:

```
symppath -sid <SymmID> disable -initiator_registration
```

- To enable PowerPath host registration on an array use:

```
symppath -sid <SymmID> enable -host_registration
```

- To disable PowerPath host registration on an array use:

```
symppath -sid <SymmID> disable -host_registration
```

Introduce devices to a host

After reconfiguring an array by moving, deleting, adding, or modifying one or more devices, update the host so that the host recognizes the new array configuration. For some platforms, the `symcfg scan` command is available to perform the host update. These include Sun Solaris, HP-UX, IBM AIX, Tru64/OSF1, and Windows systems.

After mapping devices to a host or changing device channel addresses, the following actions are required to introduce devices for each of the following host types.

Introduce devices to Sun Solaris systems

Syntax

To add or modify devices while online in the Solaris environment, use the following form:

```
drvconfig | symcfg scan  
disks  
devlinks
```

 **NOTE:** Follow the instructions in the Solaris documentation to introduce new devices to the host environment.

Introduce devices to HP-UX systems

Syntax

To view mapping change results for HP-UX hosts, use the `ioscan` command in a statement similar to the following:

```
ioscan -fnC disk
```

To define newly connected physical volumes to the HP-UX host system without rebooting it, use the following form:

```
insf -e
```

 **NOTE:** For more information, refer to the HP 9000 documentation.

Introducing devices to IBM AIX systems

About this task

To introduce new devices for AIX hosts, perform the following actions:

Steps

1. From the **SMIT** menu, select **Devices > Fixed Disk > Add a Disk**.
2. Select the **EMC SYMMETRIX** definition from the disk table.
3. Select the SCSI bus on which the new disk resides.
4. Type the connection address for the new device (target, LUN).
5. Select **EXECUTE**.
6. Repeat steps 2 through 5 for each new device being added to the configuration.

Introducing devices to HP Tru64 UNIX systems

About this task

To introduce new devices for Tru64 UNIX hosts, perform the following actions:

Steps

1. At the prompt, type:

```
scsimgr -scan_bus bus=BUSNUM
```

2. Repeat for each LUN:
3. Write a label to the device you are defining:

```
disklabel -rw rz<lun_letter><unitID> <label>
```

4. Change the ownership on the device to a particular application:

```
chown <owner>:<group> *rz<lun letter><unitID>*
```

5. Follow the host documentation to introduce new devices to the host environment.

Introducing devices to Windows systems

About this task

To introduce new or changed devices for Windows hosts, while the system remains online:

Steps

1. From the desktop, select **Start > Settings > Control Panel > Add /Remove Hardware**. Complete the wizard to discover and add the new devices.
2. Partition and format the new devices as described in the documentation for the specific Windows OS version.

Set SRDF group attributes

Description

SRDF group attributes allow you to assign priorities to SRDF/A sessions, and to set the minimum amount of time before attempting an SRDF/A cycle switch. Use the `symrdf` command to set SRDF attributes for arrays running Solutions Enabler. Refer to the Dell EMC Solutions Enabler SRDF Family CLI User Guide for more details.

NOTE:

Starting with Solutions Enabler V8.0.1, the ability to set SRDF group attributes using the `symconfigure` command is no longer available. Use the `symrdf` command to set SRDF attributes for arrays running Solutions Enabler.

Swap RA groups

Description

Use the `symrdf` command to set swap devices in an RA group from target to source. Refer to the Dell EMC Solutions Enabler SRDF Family CLI User Guide for more details.

NOTE:

Starting with Solutions Enabler V8.0.1, the ability to swap SRDF group attributes using the `symconfigure swap ra group` command is no longer available. Since `symconfigure` support for performing application-specific functions requires more restrictive access and authorization rights, generally requires more system resources (due to the session management overhead) and is more time-consuming, this functionality is no longer supported.

Virtual Witness (vWitness)

Virtual Witness (vWitness) is an additional resiliency option introduced in HYPERMAX OS 5977.945.890 and Solutions Enabler or Unisphere V8.3. vWitness has similar capabilities to the Array Witness method, except that it is packaged to run in a virtual appliance (vApp) on a VMware ESX server, not on an array. There can be up to 32 vApps, each providing a *vWitness instance*.




Figure 7. SRDF/Metro vWitness vApp and connections

The R1 and R2 arrays each contain a user-defined list of *vWitness definitions* that identify the vWitness instances that the array can use. A vWitness definition consists of a user-specified name and the location of the instance (either the IP address or the fully-qualified DNS name). The lists of vWitness definitions on each array do not have to be identical. However, they must have at least one definition in common. Initially, the R1 and R2 arrays negotiate which vWitness instance to use from the list of vWitness definitions that each array holds.

Should the SRDF links between the R1 and R2 arrays fail, or one of the arrays has a serious problem, the vWitness instance decides which array remains available to the host or hosts.

Unisphere for VMAX and SYMCLI provide facilities to manage a vWitness configuration. The user can add, modify, remove, enable, disable, and view vWitness definitions on the arrays. In addition, the user can add and remove vWitness instances. To remove an instance, however, it must not be actively monitoring SRDF/Metro activities.

vWitness requirements

vWitness requires the following:

- Array requirements:
 - Two VMAX arrays running HYPERMAX OS 5977.945.890 or later.
 - SRDF/Metro license installed on each array.
 - eManagement guest for Unisphere on each array. eManagement is standard on VMAX All Flash arrays, and can be added to VMAX3 arrays in the field. Contact your EMC representative for more information.
 - RA (Fibre/SAN) or RE (Ethernet/IP) connectivity between the paired arrays.
 - Ethernet/IP connectivity between each array and each vWitness instance it uses.
- vApp host requirements:
 - VMware ESX 4.0 or higher
 - Depending on the vApp, the host must meet the following:

- Solution Enabler Virtual Appliance: Single processor; 2 GB of memory; dual disks, with 16 GB of disk space and 5 GB of expandable disk space
- Unisphere for VMAX: Dual core processor, 16 GB of memory, and 120 GB of disk space

vWitness Management

Solutions Enabler CLI commands are available to configure, manage, and monitor a storage system's access to vWitness instances. The CLI allows for the following vWitness operations:

- Add vWitness definition
- Enable vWitness definition
- Modify vWitness definition
- Remove vWitness definition
- Suspend vWitness definition
- View vWitness definitions

Refer to the Dell EMC SRDF/Metro vWitness Configuration Guide for Solutions Enabler command syntax and examples.

Add director

NOTE: Adding the first RF Director will also enable SRDF. From Solutions Enabler V8.4 and higher, no BIN file change is required to add the first RF instance and enable SRDF. This enhancement adds flexibility to add and remove SRDF on VMAX3 and/or VMAX All Flash arrays for NDM migrations without the need to contact Customer Support.

Syntax

To add a director, use the following syntax:

```
add dir slot_num = <director slot_number> type=<FA|FE|SE|RF|RE>;
```

Examples

```
symconfigure -sid 084 commit -cmd "add dir slot_num = 1 type=FA;"
```

```
Execute a symconfigure operation for symmetrix '000197100084' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
Establishing a configuration change session.....Established.
Processing symmetrix 000197100084
{
add dir slot_num = 1 type=FA;
}
Performing Access checks.....Allowed.
. . .
Local: COMMIT.....Done.
Created FA type director: 1f Terminating the configuration change
session.....Done.
The configuration change session has successfully completed.
```

Restrictions

The following restrictions apply to adding directors:

- Requires the following security privileges:
 - Access Type: CFGSYM
 - Authorization Rights: Storage Admin
- Addition of the following directors 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

Remove director

NOTE: Removing the last RF Director will also disable SRDF. From Solutions Enabler V8.4 and higher, no BIN file change is required to remove the last RF instance and disable SRDF. This enhancement adds flexibility to add and remove SRDF on VMAX3 and/or VMAX All Flash arrays for NDM migrations without the need to contact Customer Support.

Syntax

To remove a director, use the following syntax:

```
remove dir director_num
```

Examples

```
symconfigure -sid 084 commit -cmd "remove dir 1f;"
```

```
Execute a symconfigure operation for symmetrix '000197100084' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
  Establishing a configuration change session.....Established.
  Processing symmetrix 000197100084
  {
remove dir 1f;
  }
  Performing Access checks.....Allowed.
  . . .
  Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

Restrictions

The following restrictions apply to adding directors:

- Access Type: CFGSYM
- Authorization Rights: Storage Admin

Port to director emulation support

Only a single emulation instance of a specific type (FA, DA, RF, EF, etc.) is available per director board. If more connectivity is needed, add additional ports to an existing emulation instance. That instance uses all cores configured to it to drive the workload across all ports assigned to it.

Each director board can contain up to 32 physical ports. These physical ports can be assigned to compatible emulation instances (and are numbered from 0-31) based on the rules outlined below. In addition, directors containing a Fibre Channel emulation have 32 virtual ports (numbered 32-63), which are reserved for use by internal guests. A capability attribute on each physical port determines the set of front-end emulations to which the port may be assigned.

The following emulation types support associating (assigning) unused ports to front-end emulations and disassociating (freeing) them:

- FA
- RF

Association of ports to EF (FICON) emulations is not supported.

Associate ports to director emulations

Description

Use the `symconfigure associate port` command to associate one or more physical ports to a director emulation type in the specified director board.

Syntax

To associate a port to a director, use the following syntax:

```
symconfigure associate port <port_num>[,<port_num>. . .] to dir <dir_num>;
```

Options

dir_num

The director number of the emulation.

Example

To associate ports 2 and 4 to emulation type FA on director 7E, enter:

```
symconfigure associate port 2,4 to dir 7E ;
```

Restrictions

- The following security privileges are required to execute this operation:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- Only free ports can be associated with director emulations. To display free ports, use the CLI command `symcfg list -sid xxx -port -free`.
- Associating ports to RE and EF emulations is not supported.
- Virtual ports cannot be associated or disassociated from an emulation.
- Ports cannot be added to IM, EDS, or DA emulations. If the specified director is running one of those emulations, the operation fails with the error SYMAPI_C_DIR_IS_NOT_A_FRONT_DIR.
- If any of the specified ports is already associated with an emulation, the operation fails with the error SYMAPI_C_PORT_ALREADY_ASSOCIATED and none of the specified ports will be associated.
- If the capabilities of any of the specified ports are incompatible with the emulation running on the director, the operation fails with the error SYMAPI_C_EMULATION_PORT_MISMATCH and none of the specified ports will be associated.
- Fibre Channel ports can only be associated with FA and RF emulations.
- Associating ports to director emulation configures the ports in `Offline` state. You have to change the state of the ports to `Online` to make them usable.

Verify port status after association

Description

After successfully assigning the ports to a given emulation on a director, issue the `symcfg list -dir ALL` command to verify that the ports are assigned to the requested director.

Examples

```
symcfg list -dir ALL -sid 064
```

Sample output

```
Symmetrix ID: 000197100064 (Local)
```

S Y M M E T R I X		D I R E C T O R S			
Ident	Type	Engine	Cores	Ports	Status
IM-7A	IM	1	1	0	Online
IM-8A	IM	1	1	0	Online
ED-7B	EDS	2	2	0	Online
ED-8B	EDS	2	2	0	Online
DF-7C	DISK	3	2	4	Online
DF-8C	DISK	3	2	4	Online
FA-7E	FibreChannel	4	1	2	Online
FA-7E	FibreChannel	4	1	4	Online
FA-8E	FibreChannel	4	1	3	Online
RF-7G	RDF_BI_DIR	5	1	1	Online
RF-8G	RDF_BI_DIR	5	1	1	Online

Disassociate ports from director emulations

Description

Use the `symconfigure disassociate port` command to disassociate one or more physical ports from an emulation type in the specified director board.

Syntax

To disassociate a port from a director type, use the following syntax:

```
symconfigure disassociate port <port_num>[,<port_num>. . .] from dir <dir_num>;
```

Options

dir_num

The director number of the emulation.

Example

To disassociate ports 1 and 4 from emulation type FA on director 7E, enter:

```
symconfigure disassociate port 1,4 from dir 7E ;
```

Restrictions

- The following security privileges are required to execute this operation:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- Disassociating ports from RE and EF emulations is not supported.
- Virtual ports cannot be associated or disassociated from an emulation.
- Passthru ports cannot be disassociated from an emulation.
- Ports cannot be removed from IM, EDS, or DA emulations. If the specified director is running one of those emulations, the operation fails with the error `SYMAPI_C_DIR_IS_NOT_A_FRONT_DIR`.
- If any of the specified ports is not associated with the specified director, the operation fails with the error `SYMAPI_C_EMULATION_PORT_MISMATCH` and none of the specified ports will be disassociated.
- If the specified director is running an FA or FE emulation and one of the specified ports is in a port group, the operation fails with the error `SYMAPI_C_DIR_HAS_PORT_IN_USE`.
- If the specified director is running an SE emulation and one of the specified ports has IP interface configured on it, the operation fails with the error `SYMAPI_C_DIR_HAS_PORT_IN_USE`.
- If the specified director is running an RF emulation and one or more of the specified ports has RDF groups defined on it, the operation fails with the error `SYMAPI_C_DIR_HAS_PORT_IN_USE`.
- Ports can be disassociated from director emulations only if they are in `Offline` state.

Verify port status after disassociation

Description

After dissociating ports from a given director emulation, issue the `symcfg list -port -free -sid 064` command to list the ports currently available to be associated with a director along with their supported interface types, maximum speed, and status.

Examples

```
symcfg list -port -free -sid 064
```

Sample output

```
Symmetrix ID: 000197100064
  Slot  Port  Flags      Speed  Status
  ----  ----  -
  1      4    ...Y...   1      Powered
  1      15   ...Y..Y   1      Powered
  2      24   .Y.....   1      Powered
  2      25   .Y.....   1      Powered
  2      30   Y....Y.   1      Powered
  2      31   Y....Y.   1      Powered
```

Legend:

Flags:

(F)A : Y = Yes, . = No
F(C)OE : Y = Yes, . = No
F(I)CON : Y = Yes, . = No
(S)E : Y = Yes, . = No
(D)X : Y = Yes, . = No
(R)F : Y = Yes, . = No
R(E) : Y = Yes, . = No

Set port characteristics

Syntax

To set the port characteristics of a specified director, use either `symconfigure` or `symcfg` command.

symcfg set

```
symcfg set -fa_loop_id <0-125>
```

```
symcfg set [enable|disable] -port_flag <VSA, NonPart, ACLX, OpenVMS, ShowACLX, SoftRst, EnvSet, DisQRst, SC3, SPC2, OS2007, ARB>
```

Options:

-fa_loop_id

Use to assign FA port address, between 0 and 125.

-port_flag

The Fibre Channel director or SE port flag name.

- *ARB*: When enabled, a SCSI bus reset only occurs to the port that received the reset (not broadcast to all channels).
- *VSA*: Enabled for VolumeSetAddressing for HP-UX hosts.
- *NonPart*: When enabled, the Fibre Channel director only uses hard-assigned addressing when it initializes on the loop. Otherwise, soft-assigned addressing is used during loop initialization (the default).
- *ACLX*: When enabled, allows storage provisioning using Auto-provisioning Groups.

- *OVMS*: Enabled for an OpenVMS fibre connection.
- *ShowACLX*: Enabled/Disabled, to make the ACLX device visible or to remove visibility from the ACLX device respectively. By default all ACLX enabled ports will have the ShowACLXDevice attribute disabled.
- *DisQRst*: When enabled, a Unit Attention (UA) that is propagated from another director does not flush the queue for this device on this director. Used for hosts that do not expect the queue to be flushed on a 0629 sense (only on a Hard Reset).
- *EnvSet*: When enabled, this flag enables the environmental error reporting by the array to the host on the specific port.
- *OS2007*: HP_UX & Windows Longhorn specific setting.
- *SC3*: When enabled, the Inquiry data is altered when returned by any device on the port to report that the array supports SCSI 3 protocol. When this flag is disabled, the SCSI 2 protocol is supported.
- *SoftRst*: When enabled for a Bull/GCOS-7 host, the array port supports the SCSI Soft Reset option.
- *SPC2*: SPC-2 in inquiry data.

symconfigure set port

NOTE: When setting port attributes if the port is online, the port is taken offline (except for ShowACLX), temporarily. When setting port attributes, it is recommended that you temporarily suspend I/O activity to the effected ports during this operation.

```
symconfigure set port DirectorNum:PortNum[FlagName = enable | disable][, ...] ]
gige primary_ip_address = IPAddress primary_netmask = IPAddress,default_gateway = IPAddress,
isns_ip_address = IPAddress primary_ipv6_address = IPAddress,primary_ipv6_prefix=<0-128>,
[fa_loop_id = Integer] [hostname = HostName];
```

NOTE: This command cannot be used to set port characteristics for iSCSI physical front end ports and iSCSI target virtual ports. Use the [modify iscsi_target](#) command.

FlagName

A SCSI or fibre port flag. Possible values for the SCSI protocol flags are in [Table 16. SCSI protocol port flags](#) on page 208, and the values for the fibre protocol flags are in [Table 17. Fibre protocol port flags](#) on page 209.

NOTE: Incorrectly changing the port flags can render the storage system inaccessible. Be sure of your needs before resetting these flags.

gige

Indicates that one or more network address values are going to be specified for a front-end Gig-E director. Addresses should use the Internet standard dot notation.

primary_ip_address

The IP address for a front-end Gig-E port.

primary_netmask

The IP netmask for a front-end Gig-E port.

default_gateway

The gateway or router address for a front-end Gig-E port.

isns_ip_address

The IP address for the Internet Storage Name Service (ISNS) associated with a front-end Gig-E port.

primary_ipv6_address

The IPv6 address for the front-end Gig-E port.

primary_ipv6_prefix

The IPv6 mask prefix for a front-end Gig-E port. The value can be 0-128, indicating the number of initial bits in the subnet that are identical.

fa_loop_id

The FA director loop ID (arbitrated loop physical address). Valid values are 0 through 125. (Hard Addressing must be enabled.) Not applicable for Gig-E ports.

hostname

The 12-character hostname.

Example:

To turn on the write protect access logix (**ACLX**) for director 7E, port 0, enter:


```
symconfigure set port 7e:0 ACLX=enable;
```

SCSI protocol flags

Table 16. SCSI protocol port flags

SCSI protocol flags	Description
Auto_Busy	When enabled specifically for Unisys A-series platforms only, this flag enables the auto-busy mechanism so that the array returns a Busy to all Unisys host requests.
Avoid_Force_Negotiate	When enabled for Sequent V4.2.3 and lower, the array never initiates negotiations. Normal array behavior is to initiate negotiations after an offline-to-online transition. This is for hosts that do not handle negotiations.
Avoid_Reset_Broadcast	When enabled, a SCSI bus reset only occurs to the port that received the reset (not broadcast to all channels).
Command_Reordering	When enabled with Tag Command Queuing in use, the incoming SCSI commands become reordered to Simple Queuing. The default is enabled and should only be disabled upon a request from EMC.
Common_Serial_Number	This flag should be enabled for multipath configurations or hosts that need a unique serial number to determine which paths lead to the same device. NOTE: This flag is not supported on arrays running HYPERMAX OS 5977. Setting this flag will return a feature not supported error.
Cyl_Count_In_Name ^a	When this flag is enabled, the array with the specified port embeds the cylinder count into the product ID returned in the SCSI Inquiry command. Enabled for Pyramid only when it is desirable to embed the array support into the Pyramid kernel.
Disable_False_Disconnect ^a	When enabled for debugging, this flag prevents the port from performing a False Disconnect operation. (Default is disabled and currently, you cannot change this flag.)
Disable_Interleaved_Cmds ^a	When enabled (always), metavolume command interleaving is being supported. This allows multiple metamembers to operate at the same time on the same volume.
Disable_Mini_Q ^a	When enabled for debugging, this flag disables the use of the Mini Queue on the port. (Default is disabled and currently, you cannot change this flag.)
Disable_Q_Reset_on_UA	When enabled, a Unit Attention (UA) that is propagated from another director does not flush the queue for this device on this director. Used for hosts that do not expect the queue to be flushed on a 0629 sense (only on a Hard Reset).
Disable_Ultra ^a	When enabled, this flag disables Ultra SCSI on an Ultra capable SA port. (Default is disabled and currently, you cannot change this flag.)

Table 16. SCSI protocol port flags (continued)

SCSI protocol flags	Description
Environ_Set	When enabled, this flag enables the environmental error reporting by the array to the host on the specific port.
Linked_Commands ^a	When enabled, this flag enables support of SCSI linked commands. It allows a host to chain SCSI commands in a manner similar to mainframe Channel Command Words (CCWs). (Default is enabled, and currently, you cannot change this flag.)
PBAY_Monitor	For the Sequent platforms only to allow emulation of the Sequent PBAY. When enabled, this flag enables low-level polling of the SCSI bus in order to intercept the nonstandard SCSI operations required for a Sequent PBAY disk subsystem. Must be used for the Sequent cluster operation for the Symmetry system for Sequent V4.2.x operating systems only. Must not be used on versions higher than V4.2.x or for any NUMA-Q systems and also not used for Fibre Channel.
SCSI_3	When enabled, the Inquiry data is altered when returned by any device on the port to report that the array supports SCSI 3 protocol. When this flag is disabled, the SCSI 2 protocol is supported.
SCSI_Support	When enabled, this flag provides a stricter compliance with SCSI standards for managing device identifiers, multi-port targets, unit attention reports, and the absence of a device at LUN 0.
Set_Qerr	This flag should be enabled for SGI platforms only to flush the queue on a contingent allegiance condition (CAC). Must be used for V5.3 and V6.2 SGI operating systems and cluster environments. Not used on versions higher than V6.2.
Soft_Reset	When enabled for a Bull/GCOS-7 host, the array port supports the SCSI Soft Reset option.
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.  NOTE: Reboot the host after setting this flag.
Wide_Transfer ^a	When enabled, this flag enables SCSI Wide operation. (Default is enabled, and currently, you cannot change this flag.)

a. Not available for host-based configuration changes.

Fibre protocol flags

The following table lists the Fibre protocol flags and their descriptions.

Table 17. Fibre protocol port flags

Fibre protocol flags	Description
ACLX	When enabled, allows storage provisioning using Auto-provisioning Groups.

Table 17. Fibre protocol port flags (continued)

Fibre protocol flags	Description
Auto_Negotiate	When enabled, allows two fibre ports to <i>handshake</i> and settle on an optimal speed for data transfer.
Non_Participating	When enabled along with the Hard_Address flag, the Fibre Channel director only uses hard-assigned addressing when it initializes on the loop. Otherwise, soft-assigned addressing is used during loop initialization (the default).
OpenVMS ^a	Enabled for an OpenVMS fibre connection.
Volume_Set_Address ^a	When enabled along with the Disk_Array flag for HP-UX hosts, the volume set addressing mode is selected. VSA mode allows octal addressing.

a. A block is added to prevent OpenVMS and Volume_Set_Address Fibre protocol port flags from being set at the same time on any given port, as setting these two flags together may result in data loss. This block will be effective for all operating system levels supported.

Setting a port to show ACLX device

Devices cannot be mapped to ACLX enabled ports. To make the ACLX device visible or to remove visibility from the ACLX device, set the Show_ACLX_Device attribute on front end ports to ENABLE or DISABLE, respectively. By default all ACLX enabled ports will have the Show_ACLX_Device attribute disabled.

Example

To make the ACLX device visible, enter:

```
symconfigure -sid 230 commit -cmd "set port 5E:0 Show_ACLX_device=ENABLE;"
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
Establishing a configuration change session.....Established.
Processing symmetrix 000197100230
{
    set port 5E:0 Show_ACLX_device=ENABLE;
}
Performing Access checks.....Allowed.
. . .
Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

Report flag details

Description

To report the new SHOW_ACLX_DEVICE port attribute configured on an array port, use either the `list -fa -v` or the `list -fa -detail` command. The following is an example output reporting the port detail for an ACLX device:

 **NOTE: Environment option SYMAPI_CTRL_OF_NONVISIBLE_DEVS in the options file must be enabled (or not present in the options file) if there is no device from the local host mapped to this port.**

Example

To list port detail for an ACLX device, enter:

```
symcfg -sid 064 list -port -fa 7e -detail
```

The following is an example output reporting the port detail for an ACLX device:

Sample output

```
Symmetrix ID: 000197100064
  S Y M M E T R I X    D I R E C T O R    P O R T S
      Ident  Port  WWN                Type          Flags  Speed  Status
      -----  ---  -----  -----  -----  -----  -----
      FA-7E   0   50000972C011C918  FibreChannel  ...X.    8   Online
      FA-7E   1   50000972C011C919  FibreChannel  XX...    8   Online
Legend:
Flags:
(A)CLX Enabled           : X = True, . = False
(S)how ACLX device Enabled : X = True, . = False, - = N/A
(V)olume Set Addressing  : X = True, . = False
(P)oint to Point        : X = True, . = False
VNX (G)ateway Direct Attach : X = True, . = False
```

Manage multiple iSCSI targets

This chapter describes how to manage multiple iSCSI targets and IP addresses on SE emulation.

Topics:

- [Manage multiple iSCSI targets overview](#)
- [Report iSCSI target information](#)
- [Expanded port group management for iSCSI targets](#)
- [Report iSCSI target port information](#)

Manage multiple iSCSI targets overview

Arrays running HYPERMAX OS 5977 Q3 2015 SR or higher allow multiple iSCSI targets (IQNs) and IP addresses on SE director emulations. Solutions Enabler, using the `symconfigure` command, provides support for managing (create, modify and delete) IP interfaces and iSCSI targets, and allows attaching and detaching IP interfaces to and from the iSCSI targets. In addition, static IP routes can be added, for routing packets going out of IP interfaces configured on SE director emulations.

The following set of rules and restrictions apply to the management of IP interfaces, iSCSI targets and IP routes:

- Each SE director port supports up to 64 IP interfaces. The maximum number of IP interfaces per SE director emulation is 1000.
- On an SE director emulation, each IP interface is uniquely identified by its IP address/network_id combination. A vlan ID (0 - 4094) must also be configured on this interface, and it must be unique across all IP interfaces defined on the same physical port.
- Each IP interface can be configured with either IPv4 or IPv6 IP address, but not both.
- The maximum number of iSCSI targets per SE director emulation is 256.
- Each iSCSI target must be assigned a network ID, and can be attached to up to 8 IP interfaces on the same SE director emulation. Each of the attached IP interfaces must have the same network ID as the iSCSI target.
- An iSCSI target is uniquely identified by its IQN (iSCSI target name, ASCII string in IQN format). The IQN can be user specified, otherwise HYPERMAX OS generates a globally-unique IQN when the iSCSI target is created. An iSCSI target can also be uniquely identified by a director emulation number and virtual port number, the iSCSI target virtual port. HYPERMAX OS assigns a iSCSI virtual port (0 - 255) when each iSCSI target is created.

NOTE: iSCSI targets are used as endpoints to the iSCSI protocol; an iSCSI virtual port is different from the SE emulation physical ports.

- The maximum number of IP routes (IPv4 + IPv6) per director emulation is 1024.

iSCSI Configuration sessions

Configuration changes in `symconfigure` sessions are done in the order that they are given. If an IP Interface or iSCSI target is created during a configuration session these can be modified or used by other configuration changes that are called for later in the same session. Therefore, if an IP Interface or iSCSI target is deleted during a configuration session, attempts to modify or use these by other configuration changes that are called for later in the same session will fail.

Arrays without embedded management and only iSCSI front-end emulations

Special operating rules exist for arrays that contain only iSCSI front-end emulations (SEs) and are running without Embedded Management. On these systems, during the installation process, HYPERMAX OS creates a bootstrap iSCSI target on one of the SE emulations, an IP Interface object (with a pre-defined IP address) on one of that SE's physical ports, and maps the ACLX device to that iSCSI Target. This provides control access through that IP address to all hosts and allows the initial provisioning of control hosts and creation of masking views. HYPERMAX OS blocks all host I/Os while this bootstrap iSCSI target exists. To enable host I/Os, delete this bootstrap iSCSI target after provisioning is established. Solutions Enabler provides the following features to support the iSCSI bootstrap target:

- A filter option is added to the `symcfg list` command that displays properties of the iSCSI bootstrap target on the array.
- With the exception of detaching from its IP Interface, all operations for the iSCSI bootstrap target are disabled. It cannot be added to a port group or a masking view, and the configuration cannot be modified (rename, modify or attach to another IP Interface).

Create IP interface on SE director port

Description

Use the `create_ip_interface` command to create an IP interface on a SE director port. Only one IP address is allowed and it must be either IPv4 or IPv6. This command requires an IP address, a subnetwork prefix, a network ID, and a vlan ID.

Syntax

To create an IP interface, use the following syntax:

```
create ip_interface dir <director_num> port <port_number>,
ip_address=<IPaddress>, ip_prefix=<ip_prefix>,
network_id = <network_id>, vlanid=<vlanid>
[, mtu = <MTU>];
```

Options

director_num

SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.

port_num

Port number on a director where the IP interface is created. Valid values are 0-31.

IPaddress

For IPv4, must be specified with dotted decimal notation format. For IPv6 must be specified in colon-hexadecimal format.

ip-prefix

For IPv4 prefix value is 1-30. For IPv6 prefix value is 1-128.

network_id

Valid values are 1 - 16383.

vlanid

Valid values are 0 - 4094.

MTU

Maximum transmission unit. Valid values are 1200 - 9000. The default value is 1500.

Examples

Create command file `create_ip_interface.cmd`:

```
create ip_interface dir 7E port 2, ip_address = 111.111.111.123,
ip_prefix = 24, network_id = 10, vlanid = 10
```

To commit command file, enter:

```
symconfigure -sid 230 -file create_ip_interface.cmd commit
```

```
Execute a symconfigure operation for symmetrix '00019710230' (y/[n]) ? y
```

```
A Configuration Change operation is in progress. Please wait...
```

```
Establishing a configuration change session.....Established.
```

```
Processing symmetrix 000197100230
```

```
{
```

```
create ip_interface dir 7E port 2, ip_address = 111.111.111.123,
ip_prefix = 24, network_id = 10, vlanid = 10;
```

```

}
Performing Access checks.....Allowed.
. . .
Terminating the configuration change session.....Done.
The configuration change session has successfully completed.

```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- For specified director port, director must be a SE director emulation.
- Port must be configured to the director emulation.
- IP address syntax must be valid syntax and the IP prefix must be in the valid range.
- Network ID must be in the valid range.
- Vlan ID must be in the valid range.
- MTU must be in the valid range.
- Maximum IP interfaces per SE director emulation is 1000.
- Maximum IP interface per SE director physical port is 64.
- Vlan ID must be unique to the IP Interface for the specified SE director physical port.
- Subnet mask of the IP Interface must be unique on a network ID within the specified SE director emulation.
- IP address must be unique within the same network ID for the specified SE director emulation.

Modify IP interface on SE director emulation

Syntax

To modify an IP interface, use the following syntax:

```

modify ip_interface dir <director_num>,
ip_address = <IPaddress>, network_id = <network_id>
[, new_network_id = <network_id>]
[, new_ip_address = <IPaddress>]
[, ip_prefix = <ip_prefix>]
[, mtu = <mtu>];

```

Options

director_num

SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.

IPaddress

For IPv4, must be specified with dotted decimal notation format. For IPv6 must be specified in colon-hexadecimal format.

ip-prefix

For IPv4 prefix value is 1-30. For IPv6 prefix value is 1-128.

network_id

Valid values are 1 - 16383.

vlanid

Valid values are 0 - 4094.

MTU

Maximum transmission unit. Valid values are 1200 - 9000. The default value is 1500.

Examples

Create command file `modify_ip_interface.cmd`:

```
modify_ip_interface dir 7E, ip_address= 111.111.111.123, network_id = 10, ip_prefix = 24;
```

To commit command file, enter:

```
symconfigure -sid 230 -file modify_ip_interface.cmd commit
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...

Establishing a configuration change session.....Established.
Processing symmetrix 000197100230
{
  modify_ip_interface dir 7E , ip_address= 111.111.111.123,
    network_id = 10, ip_prefix = 24;
}

Performing Access checks.....Allowed.
. . .
Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- Director must be a SE director emulation.
- Port must be configured to the director emulation.
- IP address syntax must be valid syntax and the ip prefix must be in the valid range.
- Network ID must be in the valid range.
- MTU must be in the valid range.
- IP address and network ID combination must exist for the director emulation.
- IP address or ip prefix must result in a unique subnet mask across the array.
- IP interface cannot be attached to an iSCSI target.
- IP address or network ID must result in a unique IP address/network ID combination across the array.
- The IP address and network_id together must be globally unique across the entire array.
- Command will fail if all properties are the same as those being changed.

Delete IP interface on SE director emulation

Syntax

To delete an IP interface, use the following syntax:

```
delete_ip_interface dir <director_num>, ip_address = <IPaddress>, network_id = <network_id>;
```

Options

IPaddress

For IPv4, must be specified with dotted decimal notation format. For IPv6 must be specified in colon-hexadecimal format.

director_num

SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.

network_id

Valid values are 1 - 16383.

Examples

Create command file `delete_ip_interface.cmd`:

```
delete_ip_interface dir 7E, ip_address= 111.111.111.123, network_id = 10;
```

To commit command file, enter:

```
symconfigure -sid 230 -file delete_ip_interface.cmd commit
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
Establishing a configuration change session.....Established.
Processing symmetrix 000197100230
{
  delete_ip_interface dir 7E, ip_address= 111.111.111.123,
  network_id = 10;
}
Performing Access checks.....Allowed.
. . .
Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- Director must be a SE director emulation.
- IP address syntax must be valid syntax.
- IP address must exist for the director emulation.
- Network ID must be in the valid range.
- IP interface cannot be attached to an iSCSI target.

Create iSCSI target on SE director emulation

Description

Use the `create_iscsi_tgt` command to create an iSCSI target on a SE director emulation. A network ID must be specified, and the IQN (iSCSI qualified name) can either be specified, or if not specified, it is auto-generated by the array. Newly created iSCSI targets are automatically assigned an iSCSI virtual port number (0-255), local to the SE director emulation, and are used to uniquely identify the iSCSI target. These virtual ports are used as endpoints to the iSCSI protocol, and are different from physical ports that are mapped to a SE emulation.

Optional settings include IP address, SCSI flags, and the TCP port.

Syntax

To create iSCSI target, use the following syntax:

```
create_iscsi_tgt dir <director_num>,
  network_id = <network_id>
  set_default_flags = <ENABLE | DISABLE>]
[, iqn = <IQN>] [, ip_address = <IPaddress> [,...]]
[, flag_name = ENABLE [,...]]
[, tcp_port=<tcp_port>];
```


Options

director_num

SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.

network_id

Valid values are 1 - 16383.

set_default_flags

Specifies whether to use iSCSI target default flags (flags listed below).

IQN

iSCSI qualified name. Must start with either "iqn." or "eui." strings and must include alphanumeric characters, colons, dashes, and periods. Maximum length is 255 characters.

IPaddress

For IPv4, must be specified with dotted decimal notation format. For IPv6 must be specified in colon-hexadecimal format.

tcp_port

Valid values are 1 - 65535. Default value is 3260.

flag_name

Valid flag values are:

- SOFT_RESET
- ENVIRON_SET
- DISABLE_q_RESET_ON_UA
- AVOID_RESET_BROADCAST
- SCSI_3 (DEFAULT is ENABLED)
- SPC2_PROTOCOL_VERSION (DEFAULT is ENABLED)
- ISID_PROTECTED
- SCSI_SUPPORT1 (DEFAULT is ENABLED)
- Volume_Set_Addressng
- OpenVMS

Examples

Create command file `iscsi_tgt_create.cmd`, with IQN specified:

```
create iscsi_tgt dir 7E, iqn = iqn.2013-06.com.emc:sn.11111111,  
network_id = 10, ip_address= 111.111.111.123;
```

To commit command file, enter:

```
symconfigure -sid 230 -file iscsi_tgt_create.cmd commit
```

```
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
```

```
A Configuration Change operation is in progress. Please wait...
```

```
Establishing a configuration change session.....Established.  
Processing symmetrix 000197100230
```

```
{  
  create iscsi_tgt dir 7E, iqn = iqn.2013-06.com.emc:sn.11111111,  
    network_id = 10, ip_address= 111.111.111.123;  
}
```

```
Performing Access checks.....Allowed.
```

```
...
```

```
Created IQN : iqn.2013-06.com.emc:sn.11111111
```

```
Committing configuration changes.....Committed.  
Terminating the configuration change session.....Done.
```

```
The configuration change session has successfully completed.
```

Create command file `iscsi_tgt_create.cmd`, without IQN specified:

```
create iscsi_tgt dir 7E, network_id = 10, ip_address= 111.111.111.123;
```

NOTE: IQN is auto-generated when the command file is committed.

To commit command file, enter:

```
symconfigure -sid 230 -file iscsi_tgt_create.cmd commit
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...

Establishing a configuration change session.....Established.
Processing symmetrix 000197100230
{
  create iscsi_tgt dir 7E, network_id = 10,
    ip_address= 111.111.111.123;
}

Performing Access checks.....Allowed.
.
.
Created IQN : iqn.1992-04.com.emc:600009700bbf824907fd019f00000001

Committing configuration changes.....Committed.
Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- Director must be a SE director emulation.
- Maximum iSCSI targets per SE director emulation is 1000.
- IP address syntax must be valid syntax.
- IQN must be unique to the array.
- IP address must exist in the director emulation.
- Network ID of the IP interface must match the network ID of the iSCSI target.
- Network ID must be a valid value.
- Maximum IP addresses per iSCSI target is 8.
- ACLX flag cannot be set or modified on an iSCSI target. This flag is always enabled.
- SHOW_ACLX_DEVICE flag cannot be set or modified on an iSCSI target. This flag is always enabled for bootstrap iSCSI targets and always disabled for user-created iSCSI targets.

Modify iSCSI target

Description

Use the `modify iscsi_tgt` command to modify scsi port attributes and the TCP port value of the iSCSI target.

Syntax

To modify an iSCSI target, use the following syntax:

```
modify iscsi_tgt
<[iqn = <IQN>] | [iscsi_dirport = <director_num>:<port_number>]>
[, flag_name=ENABLE | DISABLE [,...]][, tcp_port=<tcp_port>]
[, network_id=<network_id>];
```

Options

director_num

SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.

port_num

Port number on a director where the IP interface is created. Valid values are 0-31.

IQN

iSCSI qualified name. Must start with either "iqn." or "eui." strings and must include alphanumeric characters, colons, dashes, and periods. Maximum length is 255 characters.

network_id

Valid values are 1 - 16383.

tcp_port

Valid values are 1 - 65535. Default value is 3260.

flag_name

Valid flag values are:

- SOFT_RESET
- ENVIRON_SET
- DISABLE_Q_RESET_ON_UA
- AVOID_RESET_BROADCAST
- SCSI_3 (DEFAULT is ENABLED)
- SPC2_PROTOCOL_VERSION (DEFAULT is ENABLED)
- ISID_PROTECTED
- SCSI_SUPPORT1 (DEFAULT is ENABLED)
- Volume_Set_Addresssing
- OpenVMS

Examples

To enable `scsi_3` flag, create command file `iscsi_tgt_mod.cmd`:

```
modify iscsi_tgt, iqn = iqn.2013-06.com.emc:sn.11111111, scsi_3= enable;
```

To commit command file, enter:

```
symconfigure -sid 230 -file iscsi_tgt_mod.cmd commit
```

```
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
```

```
A Configuration Change operation is in progress. Please wait...
```

```
Establishing a configuration change session.....Established.  
Processing symmetrix 000197100230
```

```
{  
    modify iscsi_tgt, iqn = iqn.2013-06.com.emc:sn.11111111,  
    scsi_3= enable;  
}
```

```
Performing Access checks.....Allowed.
```

```
.....  
Terminating the configuration change session.....Done.
```

```
The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin

- For specified director port, director must be a SE director emulation.
- iSCSI target must exist but cannot be the bootstrap target.
- If modifying SCSI flags, iSCSI target must be in an offline state.
- Specified SCSI flag must be a valid flag.
- Network ID must be a valid value.
- Network ID of the IP interface must match the network ID of the iSCSI target.
- If modifying network ID, iSCSI target must be in an offline state.
- iSCSI target cannot be attached to an IP interface.
- ACLX flag cannot be set or modified on an iSCSI target. This flag is always enabled.
- SHOW_ACLX_DEVICE flag cannot be set or modified on an iSCSI target. This flag is always enabled for bootstrap iSCSI targets and always disabled for user-created iSCSI targets.

Delete iSCSI target

Description

To delete an iSCSI target either specify either the IQN (iSCSI qualified name) or the iSCSI virtual port equivalent as the iSCSI target.

Syntax

To delete an iSCSI target, use the following syntax:

```
delete iscsi_tgt <[iqn = <IQN>] | [iscsi_dirport = <director_num>:<port_number>]>;
```

Options

- IQN**
iSCSI qualified name. Must start with either "iqn." or "eui." strings and must include alphanumeric characters, colons, dashes, and periods. Maximum length is 255 characters.
- director_num**
SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.
- port_num**
Port number on a director where the IP interface is created. Valid values are 0-31.

Examples

Create command file `iscsi_tgt_delete.cmd`:

```
delete iscsi_tgt iscsi_dirport = 7E:0 ;
```

To commit command file, enter:

```
symconfigure -sid 230 -file iscsi_tgt_delete.cmd commit

Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y

A Configuration Change operation is in progress. Please wait...

Establishing a configuration change session.....Established.
Processing symmetrix 000197100230
{
  delete iscsi_tgt, iqn = iqn.2013-06.com.emc:sn.11111111;
}

Performing Access checks.....Allowed.
.
.
Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- For specified director port, director must be a SE director emulation.
- iSCSI target must exist.
- The iSCSI target cannot be the bootstrap target if no masking views exist on the array.
- iSCSI target must be in an Offline state.
- iSCSI target cannot be in a port group.

Rename iSCSI target

Description

To rename an iSCSI target specify either the IQN (iSCSI qualified name) or the iSCSI virtual port equivalent as the iSCSI target, and specify the new IQN.

Syntax

To rename an iSCSI target, use the following syntax:

```
rename iscsi_tgt <[iqn = <IQN>] | [iscsi_dirport = <director_num>:<port_number>]>
to new_iqn = <IQN>;
```

Options

- IQN**
- iSCSI qualified name. Must start with either "iqn." or "eui." strings and must include alphanumeric characters, colons, dashes, and periods. Maximum length is 255 characters.
- director_num**
- SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation
- port_num**
- Port number on a director where the IP interface is created. Valid values are 0-31.

Examples

Create command file `iscsi_tgt_rename.cmd`:

```
rename iscsi_tgt iqn = iqn.2013-06.com.emc:sn.11111111,
to new_iqn = iqn.2013-12.com.emc:sn.11111111
```

To commit command file, enter:

```
symconfigure -sid 230 -file iscsi_tgt_rename.cmd commit
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
Establishing a configuration change session.....Established.
Processing symmetrix 000197100230
{
  rename iscsi_tgt iqn = iqn.2013-06.com.emc:sn.11111111,
    to new_iqn = iqn.2013-12.com.emc:sn.11111111 ;
}
Performing Access checks.....Allowed.
. . .
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- For specified director port, director port must be a SE director emulation.
- Existing and new IQN naming rules apply. See [Options](#).
- iSCSI target must exist.
- The iSCSI target cannot be the bootstrap target.
- New IQN must be unique.

Attach IP interface to iSCSI target

Description

To attach an IP interface to an iSCSI target, specify either the IQN (iSCSI qualified name) or the iSCSI virtual port equivalent as the iSCSI target.

Syntax

To attach an IP address, use the following syntax:

```
attach ip_interface ip_address = <IPAddress>, to iscsi_tgt  
<[iqn = <IQN>] | [iscsi_dirport = <director_num>:<port_number>]>;
```

Options

IPAddress

For IPv4, must be specified with dotted decimal notation format. For IPv6 must be specified in colon-hexadecimal format.

IQN

iSCSI qualified name. Must start with either "iqn." or "eui." strings and must include alphanumeric characters, colons, dashes, and periods. Maximum length is 255 characters.

director_num

SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.

port_num

Port number on a director where the IP interface is created. Valid values are 0-31.

Examples

Create command file `iscsi_tgt_attach.cmd` using the IQN:

```
attach ip_interface ip_address = 10.10.10.1, to iqn = iqn.2013-06.com.emc:sn.11111111
```

Create command file `iscsi_tgt_attach.cmd` using the iSCSI virtual port:

```
attach ip_interface ip_address = 10.10.10.1, to iscsi_dirport = 7E:0;
```

To commit file, enter:

```
symconfigure -sid 230 -file iscsi_tgt_attach.cmd commit
```

```
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
```

```
A Configuration Change operation is in progress. Please wait...

Establishing a configuration change session.....Established.
Processing symmetrix 000197100230
{
  attach ip_interface ip_address = 10.10.10.1, to
    iqn = iqn.2013-06.com.emc:sn.11111111;
}

Performing Access checks.....Allowed.
. . .
Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- For specified director port, director port must be a SE director emulation.
- iSCSI target must exist.
- IP address must exist on the director emulation of the iSCSI target's with the same network ID.
- IP address must have no assignment to any iSCSI target on the director emulation.

Detach IP interface from iSCSI target

Description

To detach an IP interface from an iSCSI target, specify either the IQN (iSCSI qualified name) or the iSCSI virtual port equivalent of the iSCSI target.

Syntax

To detach an IP address, use the following syntax:

```
detach ip_interface ip_address = <IPaddress> from iscsi_tgt
  <[iqn = <IQN>] | [iscsi_dirport = <director_num>:<port_number>]>;
```

Options

IPaddress

For IPv4, must be specified with dotted decimal notation format. For IPv6 must be specified in colon-hexadecimal format.

IQN

iSCSI qualified name. Must start with either "iqn." or "eui." strings and must include alphanumeric characters, colons, dashes, and periods. Maximum length is 255 characters.

director_num

SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.

port_num

Port number on a director where the IP interface is created. Valid values are 0-31.

Examples

Create command file `iscsi_tgt_detach.cmd` using the IQN:

```
detach ip_interface ip_address = 10.10.10.1, from iscsi_tgt iqn =
iqn.2013-06.com.emc:sn.11111111 ;
```

Create command file `iscsi_tgt_detach.cmd` using the iSCSI virtual port:

```
detach ip_interface ip_address = 10.10.10.1, from iscsi_tgt iscsi_dirport = 7E:0;
```

To commit file, enter:

```
symconfigure -sid 230 -file iscsi_tgt_detach.cmd commit

Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y

A Configuration Change operation is in progress. Please wait...

Establishing a configuration change session.....Established.
Processing symmetrix 000197100230
{
  detach ip_interface ip_address = 10.10.10.1, from
    iscsi_tgt iqn = iqn.2013-06.com.emc:sn.111111111;
}

Performing Access checks.....Allowed.
.
Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- For specified director port, director port must be a SE director emulation.
- iSCSI target must exist.
- IP address must be attached to the iSCSI target.

Add IP route to SE director emulation

Syntax

To add IP route, use the following syntax:

```
add ip_route dir <director_num>,
  ip_address = <IPaddress>, ip_prefix = <ip_prefix>,
  gateway = <IPaddress>
  [, network_id = <network_id>];
```

Options

director_num

SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.

IPaddress

For IPv4, must be specified with dotted decimal notation format. For IPv6 must be specified in colon-hexadecimal format.

ip-prefix

For IPv4 prefix length is 1-32 characters. For IPv6 length is 1-128 characters.

network_id

Valid values are 1 - 16383.

Examples

Create command file `add_ip_route.cmd`:

```
add ip_route dir 7E, ip_address = 10.10.10.0, ip_prefix = 24, gateway = 10.10.9.1, network_id = 10;
```

To commit command file, enter:

```
symconfigure -sid 230 -file add_ip_route.cmd commit
```

```
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
```

```
A Configuration Change operation is in progress. Please wait...
```

```
Establishing a configuration change session.....Established.
```

```
Processing symmetrix 000197100230
```

```
{
  add ip_route dir 7E,ip_address = 10.10.10.0, ip_prefix = 24,
  gateway = 10.10.9.1 , network_id = 10;
}
```

```
Performing Access checks.....Allowed.
```

```
. . .
```

```
Terminating the configuration change session.....Done.
```

```
The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- For specified director port, director port must be a SE director emulation.
- IPv4 and IPv6 routes can be added separately, but not in the same command.
- A default route through a gateway can be specified as follows:
 - For IPv4 — `ip_address` (destination) 0.0.0.0 with `ip_prefix = 0`.
 - For IPv6 — `ip_address` (destination) ::, with `ip_prefix = 0`.
- Only one default gateway per director emulation is allowed.
- Maximum IP routes per SE director board is 1024.

Remove IP route from SE director emulation

Syntax

To remove IP route, use the following syntax:

```
remove ip_route dir <director_num>,
  ip_address = <Ipaddress>
  [, network_id = <network_id>];
```

Options

director_num

SE director where the IP interface is created. Valid values are 1 - 128 and must represent an SE director emulation.

IPaddress

For IPv4, must be specified with dotted decimal notation format. For IPv6 must be specified in colon-hexadecimal format.

network_id

Valid values are 1 - 16383.

Examples

Create command file `remove_ip_route.cmd`:

```
remove ip_route dir 7E, ip_address = 1:1:2:: , network_id = 10;
```

To commit command file, enter:

```
symconfigure -sid 230 -file remove_ip_route.cmd commit
```

```
Execute a symconfigure operation for symmetrix '000197100230' (y/[n]) ? y
```

```
A Configuration Change operation is in progress. Please wait...
```

```
Establishing a configuration change session.....Established.
```

```
Processing symmetrix 000197100230
```

```
{
  remove ip_route dir 7E, ip_address = 1:1:2:: , network_id = 10;
}
```

```
Performing Access checks.....Allowed.
```

```
. . .
Terminating the configuration change session.....Done.
```

```
The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute this command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- For specified director port, director port must be a SE director emulation.
- The IP route must exist.

Example iSCSI configuration

Configuring a Single IP interface and iSCSI target

The following example configures an iSCSI target on an SE director that can be used for provisioning LUNs to an IPv4 accessible host:

- Create an iSCSI target with a given IQN and `network_id` of 10 on an SE director emulation.
- Create an IP interface on a physical port associated to the same SE director emulation by giving it an IPv4 address and network prefix of 13.17.253.21/24, a `vlan ID` of 0 (must be unique to the port), and the `network_id` of 10.
- Attach to the iSCSI target the configured IP interface using its IPv4 address.
- Add a default IPv4 route (0.0.0.0) on the same SE director using the configured IPv4 interface for the outgoing IPv4 packets.

Configuring a Second IP interface to the same iSCSI Target

The following example configures a second IP interface to the existing iSCSI target (configured above) on the same SE director:

- Create a second IP interface on a physical port associated to the same SE director emulation by giving it a different IPv4 address and network prefix of 13.17.255.23/24, a `vlan ID` of 1 (must be unique to the port), and the same `network_id` of 10.
- Attach to the iSCSI target the configured IP interface using its IPv4 address.
- Add a default IPv4 route (0.0.0.0) on the same SE director using the configured IPv4 interface for the outgoing IPv4 packets.

The iSCSI target can now be used to make devices visible to hosts by adding it to a port group.

Symconfigure command restrictions for iSCSI targets

The following is a list of `symconfigure` command behavior for iSCSI targets:

- `symconfigure set port` — Does not support setting port characteristics for iSCSI physical front end ports and iSCSI target virtual ports. Use `modify iscsi_tgt` command to set flags on iSCSI target virtual ports.
- `symconfigure set port copying port` — Doesnot support copying port characteristics for iSCSI physical front end ports and iSCSI target virtual ports.
- `symconfigure associate port` — Supports association of one or more free iSCSI physical ports to SE director emulation. The following restrictions for this command are:
 - If any of the specified ports is already associated with an emulation, the operation fails and displays error.
 - If any of the specified ports are non iSCSI physical ports and the director emulation specified is SE, the operation fails and displays error.
 - If the specified ports are iSCSI physical ports and the director emulation specified is anything other than SE, the operation fails and displays error.
 - Association of ports to director emulation configures the ports in an *Offline* state. Ports states must be explicitly changed to *Online* to make them usable.
- `symconfigure disassociate port` — Supports disassociation of one or more iSCSI physical ports from SE director emulation. The following restrictions for this command are:
 - If any of the specified ports is not associated with the specified director, the operation fails and displays error.
 - Ports must be in *Offline* state.
 - Ports cannot have any devices mapped to it.
 - If the specified director is running an SE emulation and one of the specified iSCSI physical port is in a port group, the operation fails and displays error.
- `symconfigure convert dir` — Does not support conversion of SE director to a different director emulation type and vice versa.
- `symconfigure preview` — Can only provide the best effort checks described in the symconfigure IP and iSCSI target management operations.

Set online or offline state for iSCSI targets

Description

Some iSCSI target operations require the associated port to be an *Offline* state. Once operations are complete the port must be returned to the *Online* state to connect with a host.

Syntax

To set iSCSI target online/offline states, use the following syntax:

```
symcfg -sid <SymmID> [-noprompt] [-v]
  -SE <#> <-p <#> | -iscsi_port <#>>
    online
    offline
```

Report iSCSI target information

iSCSI target information is reported using the `symcfg list -ip`, `list _iscsi_tgt`, and `list -route` commands.

List array IP addresses

Syntax

To list the IP addresses configured on an array, use the following syntax:

```
symcfg [-sid <SymmID>] [-offline]
  list -ip [-SE <#|ALL>] [-p <#>] [-by_ip]
  list -ip [-RE <#|ALL>] [-p <#>] [-by_ip]
```

Examples

To list IP addresses configured on SE port 10H of array 230, enter:

```
symcfg -sid 230 list -ip -se 10H
```

Sample output

```
Symmetrix ID: 000197100230 (Local)
Dir:P  NetId  Vlan  IP Address  Mtu  iSCSI Port
-----
10H:00  1      1    10.10.10.1/24  1500  0
        1      2    10.10.11.1/24  1500  -
        2      3    3ffe:80c0:22c0:1049:1260:1600:3ce0:1234/64  1500  1
10H:02  1      1    11.11.11.1/24  1500  0
        4      2    11.11.12.1/24  1500  128
```

List array iSCSI targets

Syntax

To list iSCSI target configured on an array, use the following syntax:

```
symcfg [-sid <SymmID>] [-offline]
  list <-SE <#|ALL>>
      <-iscsi_tgt [-iqn <TargetIQN> | -iscsi_port <#> | -bootstrap]
          [-by_iqn] [-detail]>
```

Options

-iqn *TargetIQN*

Lists only the specified IQN.

-iscsi port #

Lists only the targets configured on the iscsi virtual port.

-bootstrap

Lists only the bootstrap iSCSI target.

-by_iqn

Display sort order is by IQN.

-detail

Lists details of IP interfaces and scsi port flags settings for each iSCSI target.

Examples

To list the iSCSI targets on SE director 10H for array 230, enter:

```
symcfg -sid 230 list -iscsi_tgt -SE 10H
```

Sample output

NOTE: Use the `SYMCLI_FULL_NAME` option to display more than 57 characters for the IQN.

```
Symmetrix ID: 000197100230 (Local)
```

```
Dir:P  NetId  Status  IQN
-----
```

```
10H:000      1 Online  iqn.1992-04.com.emc:5000097300092190
...
10H:255      4 Offline iqn.1992-04.com.emc:5000097300092194
```

List array IP routes

Syntax

To list IP routes configured on an array, use the following syntax:

```
symcfg [-sid <SymmID>] [-offline]
list -route [-SE <#>] [ -ipv4 | -ipv6 ] [-v]
```

Options

-ipv4 | -ipv6

Lists only ipv4 or ipv6 routes

Examples

To list -ipv4 route for SE director 10H on array 230, enter:

```
symcfg -sid 230 list -route -SE 10H -ipv4
```

Sample output

```
Symmetrix ID: 000197100230 (Local)
Director Identification: SE-10H
```

Network Id : 5

Destination	Gateway
87.215.195.176/29	0.0.0.0
62.58.50.0/25	62.58.236.234

Network Id : 6

Destination	Gateway
192.168.1.0/24	0.0.0.0
10.0.0.0/8	0.0.0.0

i **NOTE:** In above table 87.215.195.176/29 means destination ipv4 address is 87.215.195.176 and prefix length = 29 (equivalent to netmask = 255.255.255.248).

Output for -ipv6 route:

```
Symmetrix ID: 000197100230 (Local)
```

```
Director Identification: SE-10H
```

Network Id : 7

Destination	Gateway
::1/128	::1
8000:1:1::/64	4000:1:1::1

Expanded symcfg and sympd reporting for iSCSI targets

Description

The `symcfg list -address` and `sympd list` commands include the `-iscsi_port #` option to support iSCSI targets.

Syntax

For the `symcfg list -address` command use the following syntax with the `-iscsi_port #` option:

```
symcfg [-sid <SymmID>] [-offline] list [-SE <# | ALL>] [<-v |  
-port [-detail] [-p <#>]] [-address [-available]] [-iscsi_port <#>]
```

For the `sympd list` command use the following syntax with the `-iscsi_port #` option:

```
sympd [-offline] [-sid <SymmID>] [-v]  
list [-SA <#|ALL>] [-p <#>] [-scsi] [-fibre]  
[-escon] [-ficon] [-gige | [-iscsi_port <#>]]  
[-powerpath] [-vcm | -aclx] [-pdevfile] [-cyl]
```

Expanded port group management for iSCSI targets

Description

For port group management, the `symaccess` command includes two options to support iSCSI targets, `-iscsi_dirport` and `-iqn` options. These options are used with the following `symaccess` port group operations:

- create, delete, rename add, remove
- list and show
- set, enable, disable, delete, and list CHAP

Additional `symaccess` commands that support iSCSI targets but do require any CLI changes are:

- The `symaccess copy -type port` and `copy view` commands that copy a port group with iSCSI ports to a target array only if the target array is running HYPERMAX OS 5977 or higher.
- The `symaccess backup` and `symaccess restore` commands backup and restore port groups with iSCSI targets. These commands also backup and restore a provision view masking record with iSCSI virtual ports.

Syntax

For port group management (create, add, remove, list, set CHAP, enable CHAP and display for iSCSI targets, use the following syntax with the `-iscsi_dirport` and `-iqn` options.

```
symaccess -sid <SymmID> -name <GroupName> -type port  
create  
create -dirport <Dir>:<Port>[,<Dir>:<Port>...]  
create -iscsi_dirport <Dir>:<Port>[,<Dir>:<Port>...]  
create -iqn <TargetIQN>[,<TargetIQN>...]  
delete [-force] [-noprompt]  
rename -new_name <NewGroupName>
```

```
symaccess -sid <SymmID> -name <GroupName> -type port  
[-celerra] [-rp] [-ckd]  
add -dirport <Dir>:<Port>[,<Dir>:<Port>...]
```

```
add -iscsi_dirport <Dir>:<Port>[,<Dir>:<Port>...]
add -iqn <TargetIQN>[,<TargetIQN>...]
```

```
symaccess -sid <SymmID> -name <GroupName> -type port
[-ckd][-force][-unmap [-celerra][-rp]]
remove -dirport <Dir>:<Port>[,<Dir>:<Port>...]
remove -iscsi_dirport <Dir>:<Port>[,<Dir>:<Port>...]
remove -iqn <TargetIQN>[,<TargetIQN>...]
```

```
symaccess -sid <SymmID> | -file <backup_filename>
list -type port [-name <GroupName>] [-detail | -v]
[-dirport <Dir>:<Port> |
-iscsi_dirport <Dir>:<Port> |
-iqn <TargetIQN>]
show <GroupName> -type port
```

```
symaccess -sid <SymmID>
<-dirport <Dir>:<Port> |
-iscsi_dirport <Dir>:<Port> |
-iqn <TargetIQN>]
set chap -cred <Credential> -secret <Secret>
```

```
symaccess -sid <SymmID>
[-dirport <Dir>:<Port> |
-iscsi_dirport <Dir>:<Port> |
-iqn <TargetIQN>]
enable chap
disable chap
delete chap
```

```
symaccess -sid <SymmID> | -file <backup_filename>
list chap [-dirport <Dir>:<Port> |
-iscsi_dirport <Dir>:<Port> |
-iqn <TargetIQN>] [-v]
```

Restrictions

- Either the `-iscsi_dirport` or `-iqn` is specified in a command, not both.
- For FA director type only the `-dirport` option is used in a command.
- For SE director type on the `-iscsi_dirport` is used in a command.
- Fibre channel ports and iSCSI targets virtual ports are not allowed in the same port group.

Report iSCSI target port information

Port information for iSCSI targets is reported, with expanded displays, using the following `symaccess` commands:

- `show -type port`
- `show view`
- `list hba`
- `list devinfo`
- `list assignment`
- `list chap`
- `list logins`

Show port group with iSCSI targets

Description

The `symaccess show -type port` command is expanded to show the iSCSI target name of the iSCSI virtual target port.

Examples

```
symaccess show host1083_ports -sid 001 -type port
```

Sample output

Displays the the iSCSI target name of the iSCSI virtual target port under Director Identification.

```
Symmetrix ID           : 000197100001
Port Group Name       : host1083_ports
Last update time     : 12:57:49 AM on Mon Apr 15,2014

Director Identification
{
  Director
  Ident Port   WWN Port Name / iSCSI Target Name
  -----
  SE-15G  000  iqn.1992-04.com.emc:sn.11111111
  SE-15G  255  iqn.1992-04.com.emc:sn.11111112
```

Show masking view with iSCSI targets

Description

The `symaccess show view` command is expanded to show the iSCSI target name of the iSCSI virtual target port.

Examples

```
symaccess show view host1083_view -sid 001
```

Sample output

Displays the the iSCSI target name of the iSCSI virtual target port under Director Identification.

```
Symmetrix ID           : 000197100001
Masking View Name     : host1083_view
Last updated at       : 12:51:37 PM on Tue May 13,2014
View last update time : 01:37:41 PM on Thu Apr 02,2015
Initiator Group Name  : hba1_2
Host Initiators
{
  ISCSI : iqn.2002-06.com.host1083 [alias: api1083/api1083]
}
Port Group Name       : host1083_ports
Director Identification
{
  Director
  Ident Port   WWN Port Name / iSCSI Target Name
  -----
  SE-15G  000  iqn.1992-04.com.emc:sn.11111111
  SE-15G  255  iqn.1992-04.com.emc:sn.11111112
```


List HBAs with iSCSI targets (HYPERMAX OS 5977 or higher)

Examples

```
symaccess list hba
```

Sample output

Displays iSCSI target virtual port in the Dir:Port column.

```
symaccess list hba

Identifier          Physical Device Path          Symmetrix ID  Dir:Port
-----
iqn.2002-06.com*   c1t50060482D52D5F02d0s23    000197100266  07G:128
```

List device information with iSCSI targets

Examples

```
symaccess list devinfo -ig my_ig -sid 001
```

Sample output

Displays iSCSI target virtual port in the Dir:Port column.

```
Symmetrix ID          : 000197100001

Initiator Group Name  : my_ig
Last update time     : 01:20:37 PM on Fri Apr 18,2014

Group last update time: 01:20:37 PM on Fri Apr 18,2014

Host Initiators
{
  ISCSI : iqn.2002-06.com.host1082 [alias: api1082/api1082]
}

Sym      Dir:Port  Physical Device Name      Host      Cap
Dev      Masking View Name
-----
0080    07G:128    Not Visible                1         2063 myview2
0081    07G:128    Not Visible                2         2063 myview2
0082    07G:128    Not Visible                4         2063 myview2
0083    07G:128    Not Visible                3         2063 myview2
-----
Total Capacity                8252
```

List device assignments with iSCSI targets

Examples

```
symaccess list assignments -dev C3:C6 -sid 001
```


List no assignments:

```
symaccess list no_assignments -dirport 15E:310 -sid 001
```

Sample output

With assignments, displays iSCSI target virtual port in the Dir:Port column.

```
Symmetrix ID          : 000197100001
Sym
Dev      Identifier      Type      Dir:Port
-----
000C3    iqn.2002-06.com*      iSCSI    SE-10G:128
000C4    iqn.2002-06.com*      iSCSI    SE-10G:128
000C5    10000000C9AE1298     FIBRE    FA-7E:000
000C6    -                      -        -
```

 **NOTE:** With no_assignments, SE directors are not displayed.

List CHAP information with iSCSI targets

Examples

```
symaccess list chap -sid 001
```

Sample output

Displays iSCSI target virtual port in the Identifier column and Director Port column, and displays the iSCSI target name of the iSCSI target virtual port.

```
Symmetrix ID          : 000197100001
Director Identification : SE-9G
Director Port         : 128
iSCSI Target Name     : iqn.1992-04.com.emc:sn.11121318
Protocol              : CHAP
Identifier            Type      State      Credential
-----
SE-9G:128            N/A      ENABLED    credential
```

List login information with iSCSI targets

Examples

```
symaccess list logins -sid 001
```

Sample output

Displays iSCSI target virtual port in the Director Port field under the SE director.

```
Symmetrix ID          : 000197100001
Director Identification : FA-5E
Director Port         : 0
Identifier            Type      User-generated      Logged On
Node Name            Port Name          FCID                In                Fabric
```

```

-----
10000000c9ae1298 Fibre 10000000c9ae1298 10000000c9ae1298 2b1b00 Yes Yes

Director Identification : SE-9G
Director Port          : 128

Identifier              Type      Node Name      User-generated      Port Name      FCID      Logged On      Fabric
-----
iqn.2002-06.com*      iSCSI  api1082              api1082              2b1b00 Yes      Yes

```

List storage group demand with iSCSI targets

Examples

```

symsg -sid 584 list -demand -by_port

```

Sample output

Displays iSCSI target virtual port in the Dir:Port column.

```

Symmetrix ID:          000197100584

Director              IO Limit              Bandwidth Limit
-----
DIR:Port      Flags      Maximum Demand      Number      Port      Maximum Demand      Number      Excess
              HD      (IO/Sec)      SGs      (MB/Sec)      (MB/Sec)      (%)      SGs      (MB/Sec)
-----
05E:000      NN      0      0      1000      0      0      0      +1000
05E:001      NN      0      0      1000      0      0      0      +1000
09G:128      NN      0      1      -      0      -      0      -

```

Output with -v (verbose) option, displays the iSCSI target virtual port in the Dir:Port column, and the iSCSI Target Name of the iSCSI target port (or virtual port) and the WWN Port Name of fibre channel ports.

```

Symmetrix ID:          000197100584

Director Identification      : FA-7E
Director Port                : 0
WWN Port Name              : 5000097300092150
Port Total Demand (IO/Sec)   : 0
Number of SGs without Limit (IO/Sec): 1

Port Negotiated Speed (MB/Sec) : N/A
Port Total Demand (MB/Sec)     : 1000
Percent Port Capability (%)    : N/A
Port Excess (MB/Sec)          : N/A
Number of SGs without Limit (MB/Sec): 0

Storage Groups (1)
{
  -----
  Name              Maximum Demand
  MB/Sec      (%)      IO/Sec
  -----
  app_psg1         1000      0      NoLimit
  app_csg1         1000      0      NoLimit
  app_csg2         NoLimit   0      NoLimit
}

. . .
Director Identification      : SE-9G
Director Port                : 128
iSCSI Target Name         : iqn.1992-04.com.emc:sn.11121318
Port Total Demand (IO/Sec)   : 0

```

```
Number of SGs without Limit (IO/Sec):      1
.....
```

Manage storage environment for VMware VVols

This chapter describes how to manage a storage environment for VVols using the Solutions Enabler CLI.

Topics:

- [Storage management for VVols overview](#)
- [Solutions Enabler CLI support for VVol management](#)
- [CLI command support for Protocol Endpoint \(PE\) devices for VVols](#)
- [Report storage containers for VVols](#)
- [Report PE and VVol devices](#)
- [Unsupported operations/features for VASA protocol endpoints](#)

Storage management for VVols overview

NOTE: VVols are supported on VMAX3 arrays VMAX All Flash arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

VMware VVols allow data replication, snapshots, encryption etc. 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. Types of VVols are:

- Config-VVol — Stores virtual machine configurations, metadata, logs, etc.
- Data-VVol — Stores operating system, application binary and user data.
- Swap-VVols — Used for memory swaps.
- Memory-VVol — Stores VVol snapshots and clones.

To support management capabilities of VVols, the storage/vCenter environment requires the following:

- EMC VMAX VASA Provider – The VASA Provider (VP) is a software plug-in that uses a set of out-of-band management APIs (VASA version 2.0). The VASA Provider exports storage array capabilities and presents them to vSphere through the VASA APIs. VVols are managed by way of vSphere through the VASA Provider APIs (create/delete) and not with the Unisphere for VMAX user interface or Solutions Enabler CLI. After VVols are setup on the array, Unisphere and Solutions Enabler only support VVol monitoring and reporting.
- Storage Containers (SC) – Storage containers are chunks of physical storage used to logically group VVols. SCs are based on the grouping of Virtual Machine Disks (VMDKs) into specific Service Levels. SC capacity is limited only by hardware capacity. At least one SC per storage system is required, but multiple SCs per array are allowed. SCs are created and managed on the array by the Storage Administrator. Unisphere and Solutions Enabler CLI support management of SCs.
- Protocol Endpoints (PE) – Protocol endpoints are the access points from the hosts to the array by the Storage Administrator. PEs are compliant with FC and replace the use of LUNs and mount points. VVols are "bound" to a PE, and the bind and unbind operations are managed through the VP APIs, not with the Solutions Enabler CLI. Existing multi-path policies and NFS topology requirements can be applied to the PE. PEs are created and managed on the array by the Storage Administrator. Unisphere and Solutions Enabler CLI support management of PEs.

Figure 8. VMAX3 VVol Architecture

VMAX3 VVol Architecture

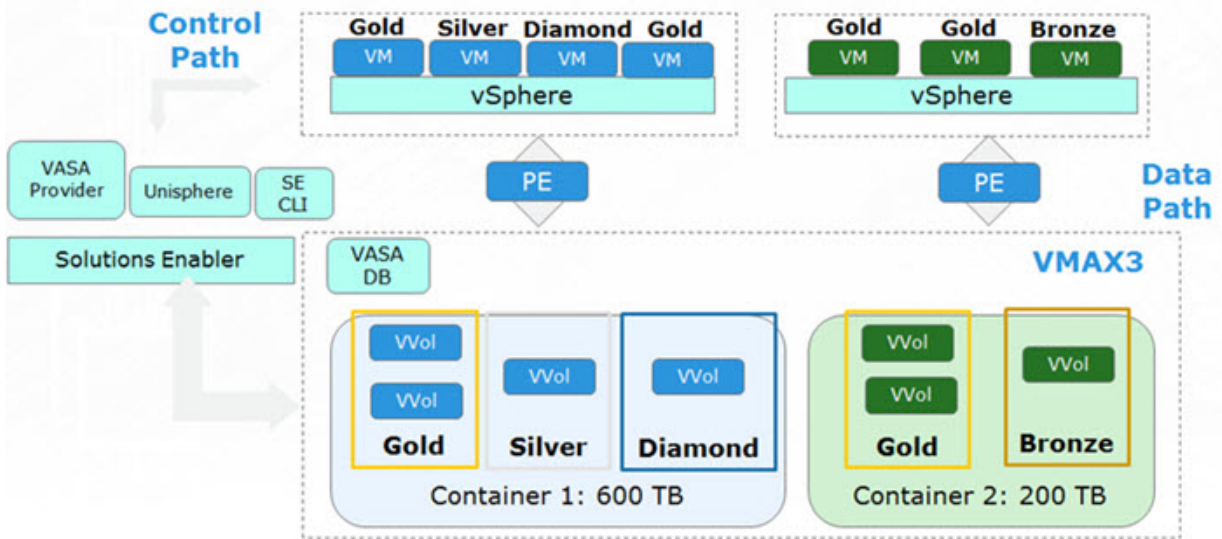


Table 18. VVol architecture component management capability

Functionality	Component
VVol device management (create, delete)	VASA Provider APIs / Solutions Enabler APIs
VVol bind management (bind, unbind)	
Protocol Endpoint device management (create, delete)	Unisphere/Solutions Enabler CLI
Protocol Endpoint-VVol reporting (list, show)	
Storage Container management (create, delete, modify)	
Storage container reporting (list, show)	

Solutions Enabler CLI support for VVol management

This section describes the Solutions Enabler CLIs modified to support VVol management and includes the following CLI actions:

- Create/delete storage containers
- Modify storage container descriptions
- Add/modify storage resource to storage containers
- Remove storage resource from storage containers
- Create Protocol Endpoint (PE) devices

For VVol management functions (create/delete VVols, PE binding) refer to the EMC VASA Provider documentation or applicable VMware vCenter documentation.

Create storage container for VVol management

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Description

Creates a new storage container on a specified array.

Syntax

To create a storage container, use the following syntax:

```
symcfg -sid <SymmID> -sc create -name <StorageContainer> -type v vols -description <Description>
```

The following security privileges are required to use this command:

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin.

Options

-name

Storage container name. Must not exceed 63 characters in length, must begin with an alpha-numeric character and may contain hyphens and underscore characters. Names are not case sensitive but is case preserving.

-description

Storage container description. Cannot exceed 128 characters, and contains only the following characters: a- z A- Z 0-9 _ ! @ # \$ % ^ & * () - . along with the space character.

-type

The only option is `v vols`.

Rules and restrictions:

Maximum number of storage containers allowed per array is 16.

Delete storage container for VVols

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Syntax

To delete a storage container, use the following syntax:

```
symcfg -sid <SymmID> -sc delete -sc_name <StorageContainer>
```

The following security privileges are required to use this command:

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin.

Rules and restrictions:

Storage container must not have any devices consuming any of the resources defined in the container.

Modify description for storage container for VVols

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Syntax

To change a storage container description, use the following syntax:

```
symcfg -sid <SymmID> -sc set -sc_name <StorageContainer> -description <Description>
```

The following security privileges are required to use this command:

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin.

Options

-description

Storage container description. Cannot exceed 128 characters, and contains only the following characters: a- z A-Z 0-9 _ ! @ # \$ % ^ & * () - . along with the space character.

Add storage resource to storage container for VVols

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Syntax

To add a storage resource to a storage container, use the following syntax:

```
symcfg -sid <SymmID> -sc -sc_name <Storage Container>  
  add -sresource <StorageResourceName>  
  -sl <SLName> -wl <WorkloadName>  
  -srp <SRPName> [-nocompression]  
  -subscribed_max <GB>
```

The following security privileges are required to use this command:

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin.

Options

-sresource

Storage resource name. Must be unique and not exceed 63 characters in length. Must begin with an alpha-numeric character and may contain hyphens and underscore characters. Names are not case sensitive but is case preserving.

-sl

Service Level name must be supplied, along with optional workload and the maximum amount of subscribed storage, in GBs, provisioned to the storage container. Service Level name must be unique to the container.

-wl

Workload name must be unique to the container.

-srp

If SRP name is not provided, then the default SRP for the FBA emulation is used.

-nocompression (-noc)

When adding a resource to a storage container the compression attribute is enabled by default on FAST managed storage groups if the associated SRP supports compression. The compression attribute is removed using this option. Compression is allowed only on VMAX All Flash Array and only FBA devices.

Rules and restrictions:

Maximum number of resources per container is 32.

Modify storage resource for storage container for VVols

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Syntax

To change the storage size limit for the storage container, use the following syntax:

```
symcfg -sid <SymmID> -sc -sc_name <StorageContainer>
set -sresource <StorageResourceName> -subscribed_max <GB>
```

The following security privileges are required to use this command:

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin.

Rules and restrictions

Subscribed limit must be the same or greater than the current aggregate subscription of devices using the resource.

Remove storage resource from storage container for VVols

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Syntax

To remove a storage resource from a storage container, use the following syntax:

```
symcfg -sid <SymmID> -sc -sc_name <StorageContainer> remove -sresource <StorageResourceName>
```

The following security privileges are required to use this command:

- Required Access type: CFGSYM

- Required Authorization Rights: Storage Admin.

Rules and restrictions

There must be no devices using the resource.

Create Protocol Endpoint (PE) devices for VVols

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Description

A PE identifies an access point to an array for one or more VMWare virtual machines, through which a VMWare Virtual Volume (VVol) receives IO. From a VM host, PEs provide a connection point for the management of very large numbers of VVols. A PE provides a data path from ESXi vSphere to VVols but does not have a backing storage.

NOTE: A PE must be in a storage group that belongs to a masking view so it is visible to the host.

Syntax

To create PE devices, use the following syntax:

```
symdev -sid <SymmID> create -pedev [-N <#>]
```

Options

-N, #

Number of PEs to create. Maximum per array is 1,024.

CLI command support for Protocol Endpoint (PE) devices for VVols

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Description

The following table lists the Solutions Enabler CLI commands that support PE devices, but do not require syntax changes for this support. Refer to *EMC Solutions Enabler Command Reference Guide* for command syntax.

Table 19. Supported CLI commands for PE devices

Command	Rules and restrictions
<code>symdev delete</code>	<ul style="list-style-type: none"> • Specified storage group cannot be FAST managed. • Storage group cannot contain VVols.
<code>symconfigure delete</code>	<ul style="list-style-type: none"> • PE cannot be bound to a VVol or group of VVols. • PE cannot belong to a storage group with VVols.
<code>symconfigure set device_name</code>	For PEs, only the <code>device_name</code> option can be used with this command.
<code>symaccess add</code> (allows user access to add PEs to storage group)	Specified storage group cannot be FAST managed.
<code>symmsg add</code> (allows adding PE devices to storage group)	Storage group cannot contain any other PEs.
<code>symmsg remove</code> (allows removing PE devices from storage group)	PE cannot be bound to a VVol.

Table 19. Supported CLI commands for PE devices (continued)

Command	Rules and restrictions
<code>symsg set</code> (there are limited set options for PE devices)	For a storage group with PE devices, command cannot request FAST management for the storage group.

Report storage containers for VVols

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Use the `symcfg list -sc` and `symcfg show -sc` commands to display storage containers for VVols. The `list` command displays the storage containers on a specified array. The `show` command displays the resources and subscription capacities of a specified storage container.

Show storage container for vVols

This topic applies only to arrays running HYPERMAX OS 5977 or higher.

Syntax

To show details and storage resources for a storage container, use the following syntax:

```
symcfg -sid <SymmID> -sc show -sc_name <StorageContainer>
```

Options

`-detail`

Shows the used subscribed capacity for the storage container.

Examples

To list the storage containers for array 0084, enter:

```
symcfg -sid 0084 show -sc -sc_name Datastore1
```

Sample output

NOTE: The storage resource names in a Storage Container display in alphanumeric order. Up to 24 characters of each storage resource, and 15 characters of Service Level and SRP name display. If the storage resource name exceeds 24 characters, the first 23 characters are displayed, followed by the "*" character to denote that the name is truncated. If the Service Level or SRP name exceeds 15 characters, the first 14 characters are displayed, followed by the "*" character to denote that the name has been truncated. If the `SYMCLI_FULL_NAME` environment variable is set none of the names are truncated.

```
Symmetrix ID          : 000197800084
Name                  : Datastore1
Type                  : VVOLS
Description           : 20TB Diamond Compressed
Subscribed Capacity Limit (GB) : 20000.0
Subscribed Capacity (GB) : 1300.1
Subscribed Capacity (%) : 7

Storage Resources (1):
{
-----
Name                Flg Service Level      SRP          Capacity
C Name              Workload Name  Limit Subs   Comp
-----
Diamond_oltprep     X   Diamond            OLTP_REP    DEFAULT_SRP 20000.0    1.5:1
-----
Total
}
                20000.0    1.5:1
```

```
Legend:
Flags:
(C)ompression      X = Compression Enabled, . = N/A
```

Output using the `-detail` option to display used subscribed capacity:

```
Symmetrix ID      : 000197800084
Name              : Datastore1
Type             : VVOLS
Description      : 20TB Diamond Compressed
Subscribed Capacity Limit(GB) : 20000.0
Subscribed Capacity (GB) : 1300.1
Subscribed Capacity (%) : 7

Storage Resources (1):
{
-----
Name              Flg Service Level      SRP          Capacity
C               Name          Workload Name  (GB)         Subs         Comp
-----
Diamond_oltprep  X   Diamond          OLTP_REP  DEFAULT_SRP  20000.0     1300.1     7   1.5:1
-----
Total              20000.0     1300.1     7   1.5:1
}

Legend:
Flags:
(C)ompression      X = Compression Enabled, . = N/A
```

List storage containers for VVols

Syntax

To list storage containers, use the following syntax:

```
symcfg -sid <SymmID> -sc list
```

Options

`-detail`

Lists the type of storage container (`vvols` is the only supported type).

`-v (-verbose)`

A "show" of all storage containers on the array are displayed.

Examples

To list the storage containers for array 0084, enter:

```
symcfg -sid 0084 list -sc -detail
```

Sample output

NOTE: Storage container names display in alphanumeric order. Up to 32 characters of each Storage Container name displays. If the name exceeds 32 characters, the first 31 characters are displayed, followed by the "*" character to denote that the name is truncated. If the `SYMCLI_FULL_NAME` environment variable is set, the storage container name is not truncated.

```
Symmetrix ID      : 000197100084

Container Name      Description
-----
Datastore1         20TB Diamond Compressed
Datastore2         10TB Diamond storage
```

Using the `-v` and `-detail` options to display container resource details:

```

Symmetrix ID           : 000197800084
Name                   : Datastore1
Type                   : VVOLS
Description            : 20TB Diamond Compressed
Subscribed Capacity Limit(GB) : 20000.0
Subscribed Capacity (GB)   : 1300.1
Subscribed Capacity (%)    : 7

Storage Resources (1):
{
-----
-----
Subs      Comp      Flg Service Level      SRP      Capacity
Name      Ratio      C  Name      Workload Name      Limit Subs
(GB) (%)
-----
-----
Diamond_oltprep      X  Diamond      OLTP_REP DEFAULT_SRP      20000.0
1300.1  7  1.5:1
-----
-----
Total
1300.1  7  1.5:1
}

Name                   : Datastore2
Type                   : VVOLS
Description            : 10TB Diamond storage
Subscribed Capacity Limit(GB) : 10000.0
Subscribed Capacity (GB)   : 1085.0
Subscribed Capacity (%)    : 1

Storage Resources (1):
{
-----
-----
Subs      Comp      Flg Service Level      SRP      Capacity
Name      Ratio      C  Name      Workload Name      Limit Subs
(GB) (%)
-----
-----
Diamond      .  Diamomd      <none>  DEFAULT_SRP      10000.0
1085.0  1  1.5:1
-----
-----
Total
1085.0  1  1.5:1
}
. . .
Legend:
Flags:
(C)ompression      X = Compression Enabled, . = N/A

```

Report PE and VVol devices

NOTE: VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Use the following commands to display PE and VVol devices:

- `symdev list`
- `symdev show`

 **NOTE:** The list commands display PE devices by default along with other device types.

List PE devices and VVol devices

Description

The `symdev list` command with no filters lists all devices including PE devices, but not VVols. Filters are used to list only PE or VVols devices.

Syntax

To list PE devices or VVol devices, use the following syntax:

```
symdev -sid <SymmID> list -vvol | -pedev | -nopdev
```

Options

-pedev

Lists only PE devices. Abbreviated `-ped`.

-vvol

Lists only VVol devices. Abbreviated `-vvo`.

-nopdev

Excludes PE devices from device list. Abbreviated `-nope`.

Examples

To list only the PE devices for array 064, enter:

```
symdev list -sid 064 -pedev
```

To list only the VVol devices for array 064, enter:

```
symdev list -sid 064 -vvol
```

Sample output

Output with no filters applied (lists PEs and other devices, VVols do not list without the `-vvol` filter):

```
Symmetrix ID: 000194900064
-----
Device Name          Dir          Device
-----
Sym  Physical        SA :P Config   Attribute  Sts  Cap
-----
. . .
00105 Not Visible    ????:? TDEV      N/Grp'd    NR   2063
00106 Not Visible    ????:? PE       N/Grp'd    RW    4
. . .
```

Output with `-vvol` filter applied:

```
Symmetrix ID: 000194900064
-----
Device Name          Dir          Device
-----
Sym  Physical        SA :P Config   Attribute  Sts  Cap
-----
. . .
```

```
00160 Not Visible      ????:? VVOL          N/Grp'd      RW      2063
. . .
```

Output with -pedev filter applied:

```
Symmetrix ID: 000194900064

      Device Name      Dir      Device
-----
Sym   Physical      SA :P Config      Attribute      Sts      Cap
-----
      (MB)
. . .
00106 Not Visible    ????:? PE          N/Grp'd      RW      4
. . .
```

Unsupported operations/features for VASA protocol endpoints

NOTE: Protocol Endpoints/VVols are supported on VMAX3 arrays running HYPERMAX OS 5977 Q1 2016SR or higher.

Feature	Description	Affected CLI
RDF (Remote Data Facility)	Pairing of a PE device with another device on a remote array is not allowed.	symrdf
TimeFinder	Replication of PE devices or any object that contains PE devices is not allowed.	<ul style="list-style-type: none"> • symmir • symclone • symsnapvx
ORS (Open Replicator)	Replication of PE devices or any object that contains PE devices is not allowed.	symrcopy
QOS (Quality of Service)	QOS of PE devices or any object that contains PE devices is not allowed.	symqos
ACL (Access Control Logic)	ACL operations of PE devices or any object that contains PE devices is not allowed.	symacl
Base control	Base control CLI does not support VASA PEs or any object containing PEs.	<ul style="list-style-type: none"> • symdev (except symdev list) • sympdev
Groups	Device group and composite group CLIs do not support VASA PEs or any object that contains PE devices. NOTE: Storage group management is supported for VASA PE. Storage group must belong to a masking view so the PE is visible to a host.	<ul style="list-style-type: none"> • symdg • sympd

Array integration with RecoverPoint

This chapter describes the RecoverPoint Integration feature and the `symrpi` command used to manage storage groups for RecoverPoint protection.

Topics:

- [RecoverPoint integration overview](#)
- [RecoverPoint integration naming conventions](#)
- [RecoverPoint device rules and restrictions](#)
- [RecoverPoint integration control operations](#)
- [RecoverPoint integration reporting](#)

RecoverPoint integration overview

The RP Integration feature integrates these arrays with an external RP appliance cluster, providing data protection using RP Continuous Data Protection (CDP) between arrays and Continuous Remote Replication (CRR) between data centers. TimeFinder SnapVX technology is the underlying replication technology used for this implementation. The integration includes a set of RP hosts configured as a cluster, and the HBAs on those hosts are zoned to an array through FA ports. Solutions Enabler, through the `symrpi` command, creates and tags RP storage groups on the array, and then by adding devices from existing storage groups on the array to the RP storage group, makes the devices visible to the RP cluster. This enables the RP system to replicate the volumes to local or remote arrays.

RecoverPoint integration support:

- VMAX3 and VMAX All Flash arrays running HYPERMAX OS 5977 Q217SR
- PowerMax arrays running PowerMaxOS

RecoverPoint integration naming conventions

Solutions Enabler must create storage groups for the explicit use by the RecoverPoint Cluster and array CLI to protect them from illegal commands. The RP integration environment requires the following object naming conventions:

- The first storage group uses the format `_RP_<ClusterName>_SG`. Subsequent storage groups must use the format: `_RP_<ClusterName>_SG<n>`.
- Cluster names
 - Maximum length of 32 characters
 - Must begin with alphanumeric character
 - Embedded hyphens are allowed, underscores are not allowed
 - Not case sensitive and are reported in upper case

RecoverPoint device rules and restrictions

This section outlines the Solutions Enabler control operation rules and restrictions for RP devices for arrays running HYPERMAX OS 5977 Q217SR or higher.

Base controls

For base controls `hold`, `unhold`, `ready`, `not_ready`, `write_disable`, `rw_enable` following rules apply:

- Base control operations are not allowed on `RP_Internal` devices for the `hold`, `unhold`, `ready` and `not_ready` actions.
- Base control operations are only allowed on `RP_Internal` devices for the `rw_enable` and `write_disable` actions when initiated by the RP system.
- Base control operations are only allowed on `RP_External` devices when initiated by the RP system.

The `symdev`, `symcg`, `symsg`, and `symsg` commands are modified to comply with these rules.

Device configuration change

Array devices used by the RP system cannot be explicitly tagged through the use of device setting commands. Devices are tagged as RP devices when an application storage group is RP protected. Solutions Enabler configuration change operations, except for valid RP actions or `symrpi` commands, are denied access to *RP_Internal* devices. The following lists the device configuration restrictions for RP-tagged devices:

- *RP_Internal* devices cannot be RDF devices.
- *RP_External* devices cannot be RDF R1, R11, R2, R22, R21, R1BCV, or RDF Metro devices.
- *RP_External* or *RP_Internal* devices cannot have the BCV attribute.

The `symconfigure` command is modified to comply with these rules.

Provisioning

A permissions flag is no longer required for the following actions on masking views (both *RP_Internal* or *RP_External* devices):

- Creating a masking view that will contain RP-tagged devices.
- Adding a device to a masking view that contains RP-tagged devices.
- Deleting a masking view that contains RP-tagged devices.
- Removing a port group from a masking view that contains RP-tagged devices.
- Removing a device from a masking view that contains RP-tagged devices.

Adding *RP_Internal* devices to a masking view, other than a RP masking view is not allowed.

The `symsg` (add/remove operations) and `symaccess` commands are modified to comply with these rules.

Open Replicator

Open Replicator operations are not allowed on RP-tagged control devices (external or internal). If the `symrcopy create` command is issued it fails and an error is returned.

TimeFinder

The following rules apply to TimeFinder operations:

- TimeFinder (Mirror, Clone, Snap) operations are not allowed on *RP_Internal* or *RP_External* devices.
- TimeFinder SnapVX operations are allowed when the source or target device is a *RP_External* device.
- TimeFinder SnapVX operations are allowed for valid RP application when the target device of the device pair is a *RP_Internal* device.

The `symsnapvx` command is modified to comply with these rules.

RDF

The following rules apply to RDF operations:

- *RP_Internal* devices cannot be RDF devices.
- *RP_External* cannot be RDF R1, R11, R2, R22, R21, R1BCVs or RDF Metro devices.
- Since R1 devices can be changed to R2 devices during a RDF swap or failover operation, the following commands are blocked, when issued to a RDF pair, where the R1 device is a *RP_External* device:
 - `symrdf swap`
 - `symrdf failover -establish`
 - `symrdf failover -restore`

The `symrdf` command is modified to comply with these rules.

RDF/Star

RP_Internal and *RP_External* devices cannot be Star devices.

The `symstar` command is modified to comply with these rules.

RecoverPoint integration control operations

The RecoverPoint integration process is managed by the `symrpi` CLI command and includes the following operations:

- Environment setup or remove
- Environment expand
- Create repositories
- Create or delete journal
- Protect or unprotect storage groups
- List any or all of the RP clusters configured on an array

RecoverPoint Environment setup - description and actions

Description

The environment setup operation creates the infrastructure that allows host-visible devices in application storage groups to be visible to the RP Appliance Cluster. Prior to running this action the storage administrator must do the following:

- Create an initiator group using the naming convention `RP_<ClusterName>_IG` and populate it with WWNs of the HBAs on the RP Appliance Cluster nodes.
- Create a port group using the naming convention `RP_<ClusterName>_PG` and add it to the array front-end ports that are zoned to the RP appliance HBAs.
- If the RP environment is expected to grow, then create multiple port groups, using different director:port pairs, when creating the groups. The naming convention for expandable port groups is `RP_<ClusterName>_PG<n>`, where `<n>` is an index between 1 and 15.

NOTE: This operation is restricted and must be run by the storage administrator on the Solutions Enabler control host. It cannot be executed by the RP system.

Actions

The environment setup operation performs the following actions:

- Validates the cluster name is unique, otherwise the operation fails.
- Validates the RP initiator group and port group, and fails the operation if they do not exist.
- If creating multiple port groups, ensures that none of the director:port pairs are the same across the port groups, otherwise the operation fails.
- Validates that each initiator in the initiator group is zoned to each port in the port group, and fails the operation if not zoned properly.
- Validates that there is at least one login entry for an initiator from the initiator group to a port from the port group, otherwise setup operation fails.
- Creates a RP storage group named `_RP_<ClusterName>_SG` and assigns an internal ID.
- Creates a RP storage group named `_RP_<ClusterName>_Attic` that is used during environment removal actions.
- Creates six gatekeeper devices that are tagged *RP_Internal*.
- Adds the six GK devices to the RP storage group.
- If the `-repository` flag is specified in the environment `-setup` command, creates a 5.7 GB repository file tagged as *RP_internal* and adds it to the RP storage group.
- Creates the masking view `RP_<ClusterName>_MV` using initiator group `RP_<ClusterName>_IG`, port group `RP_<ClusterName>_PG`, and storage group `_RP_<ClusterName>_SG`.
- If other Port Groups exist with the naming convention `RP_<ClusterName>_PG<n>`, creates the masking view `RP_<ClusterName>_MV<n>`, using initiator group `RP_<ClusterName>_IG`, port group `RP_<ClusterName>_PG<n>`, and storage group `_RP_<ClusterName>_SG<n>`.

Setup RecoverPoint environment

Description

The `symrpi environment -setup` command configures the RecoverPoint Cluster environment on the target array. Once the environment setup action completes, the operation continues with an environment expand action. This operation looks for port groups named `RP_<ClusterName>_PG<n>`, and if none are found the operation completes with success.

This operation requires the following security privileges:

- Access type – CFGSYM
- Authorization rights – Storage Admin

The command requires the following parameters:

- Target array (`-sid symID`)
- RP Cluster name (`-cluster ClusterName`)

Syntax

To setup the RP environment on a target array, use the following syntax:

```
symrpi -sid symID -cluster ClusterName environment -setup -repository -nop
```

Options

-repository

Creates a Repository device with the environment setup action.

-nop

Requests that no prompts are returned after the command is executed. The default is to prompt for command confirmation.

Examples

To configure the cluster `CorpRPCluster` on array 385 with a `-repository` device, enter:

```
symrpi -sid 395 -cluster CorpRPCluster environment -setup -repo

Execute 'Environment Setup' operation on cluster 'CorpRPCluster' (y/[n])? y

An RP 'Environment Setup' operation is
in progress for cluster 'CorpRPCluster'. Please wait...

  Analyze Configuration.....Validated.
  Setup RecoverPoint Environment.....Started.
  Setup RecoverPoint Environment.....Done.

Devices created: 0258B

The RP 'Environment Setup' operation successfully executed for
cluster 'CorpRPCluster'.

An RP 'Environment Expand' operation is
in progress for cluster 'CorpRPCluster'. Please wait...

The RP 'Environment Expand' operation successfully executed for
cluster 'CorpRPCluster'.
```

Expected command behavior when RP environment is already setup on the array:

```
symrpi -sid 124 -cluster CorpRPCluster environment -setup -repository

A RP 'Environment Setup' operation is in progress. Please wait...

The RecoverPoint Environment already exists for this Symmetrix
```

RecoverPoint journal device creation - description and actions

Description

The create journal devices operation creates a number of journal devices of a specified size and makes those devices visible to the RP Appliance.

NOTE: This operation can be run by the storage administrator on the Solutions Enabler control host or by the RP system.

Actions

The create journal operation performs the following actions:

- Validates that the cluster exists on the target array, otherwise the operation fails.
 - Validates that there is enough space in at least one of RecoverPoint storage groups to hold the new journal devices. Saves the internal identifier that the array assigned to that RecoverPoint storage group, and fails if the storage group does not exist. The operation also fails if none of the RP storage groups has the capacity to store the new journal devices.
 - Validates that there is adequate space in at least one of the collections of Recover Point Storage Groups, otherwise the operation fails.
 - If there is no storage group named `_RP_<ClusterName>_Journal` defined on the array, creates an empty storage group with this name, and assigns it the *Optimized* service level on a hybrid array, and the *Diamond* service level on an All-Flash array.
- NOTE:** Changing the storage group attributes on this storage group such as service level or compression is allowed. For this storage group compression is disabled by default.
- Creates N devices of the specified size that are tagged `RP_Internal`, and assigns a nice-name to each device.
 - Adds the created journal devices to the `_RP_<ClusterName>_Journal` storage group.

Create RecoverPoint journal devices

Description

The `symrpi create -journal` command creates a number of journal devices of a specified size and makes those devices visible to the RP Appliance. This operation requires the following security privileges:

- Access type – CREATEDV
- Authorization rights – Storage Admin

The command requires the following parameters:

- Target array (`-sid symID`)
- RP Cluster name (`-cluster ClusterName`)
- Size of each journal volume (`-cap n`). Default is Mb.
- Number of journal volumes (`-N n`)

Options

`-captype (cyl/mb/gb/tb)`

Capacity type for journal size.

Syntax

To create journal devices, use the following syntax:

```
symrpi -sid symID -cluster ClusterName create -journal -cap n -N n
```

Examples

To create two 2200 Mb journal devices for the RP cluster `CorpRPcluster`, enter:

```
symrpi -sid 124 -cluster CorpRPcluster create -journal -cap 2200 -N 2 -nop
```

```
A RecoverPoint 'create Journal' operation is  
in progress for cluster 'CorpRPcluster'. Please wait...
```

```
Analyze Configuration.....Validated.
Create Devices.....Started.
Create Devices.....Done.
```

```
Devices Created: 005B6 005B7
```

```
The RecoverPoint 'create Journal' operation successfully executed for cluster 'CorpRPcluster'.
```

RecoverPoint Environment expand - description and actions

Description

The environment expand operation expands the RP environment to accommodate additional Production, Replica, RP Internal volumes. This expansion operation addresses the limitation of the initial RP environment that can only accommodate 4096 devices, due to the array restriction that only 4096 volumes can be visible to a host from a masking view. This action is initiated by the RP appliance or the storage administrator, and results in the creation of a new masking view that creates an additional 4096 devices that are visible to the RP Cluster.

NOTE: If VSA (Volume Set Addressing) is enabled on the director port then the maximum RP expansion is 2000 devices.

During expansion the initial RP initiator group is reused, therefore virtual initiators in separate RP initiator groups are not used. Reuse of the initial initiator group requires that the storage administrator create additional port groups. The naming convention for the Port Groups must be `RP_<clustername>_PG<n>`, where `<n>` is consecutive index between 1 and 15. When issuing an expansion of the RP environment, the original initiator named `RP_<clustername>_IG` is used during the creation of the Recover Point masking view `RP_<ClusterName>_MV<n>`. If the Recover Point port groups don't exist during the expansion then the expansion command completes with no other Masking Views being created.

This action is not restricted so it can be executed by the RP system or run by the storage administrator on the Solutions Enabler control host.

Actions

The environment setup operation performs the following actions:

- Validates that the cluster exists, otherwise the operation fails.
- Validates that the port group exists, otherwise the operation fails.
- If creating multiple port groups, ensures that none of the director:port pairs are the same across the port groups, otherwise the operation fails.
- Validates that the initial initiator group exists.
- Validates that each initiator in the initiator group is zoned to each port in the port group, otherwise the operation fails.
- Validates that there is at least one entry in the Login History Table (LHT) linking an initiator from the initiator group to a port from `PG<n>`, otherwise the operation fails.
- Determines the next index `<n>` in the range (1 - 15) for the Port Group `RP_<ClusterName>_PG<n>`. This operation fails if:
 - all port groups do not exist.
 - if the initial initiator group does not exist.
 - if the corresponding `_RP_<clustername>_SG<n>` already exists.
 - if the corresponding `RP_<clustername>_MV<n>` already exists.
- Creates the RP storage group `_RP_<clustername>_SG<n>`.
- Adds one of the 6 GK devices from storage group `_RP_<ClusterName>_SG` to the storage group `_RP_<ClusterName>_SG<n>`.
- Creates the masking view `RP_<ClusterName>_MV<n>` using initiator group `RP_<ClusterName>_IG<n>`, port group `RP_<ClusterName>_PG<n>`, and storage group `_RP_<ClusterName>_SG<n>`.

Expand RecoverPoint environment

Description

The `symrpi environment -expand` command expands the RecoverPoint Cluster environment on the target array. This operation requires the following security privileges:

- Access type – CFGSYM
- Authorization rights – Storage Admin

The command requires the following parameters:

- Target array (`-sid symID`)
- RP Cluster name (`-cluster ClusterName`)

Syntax

To expand the RP environment on a target array, use the following syntax:

```
symrpi -sid symID -cluster ClusterName environment -expand -nop
```

Options

`-nop`

Requests that no prompts are returned after the command is executed. The default is to prompt for command confirmation.

Examples

To expand the RP environment for the cluster `CorpRPCluster` on array 124, enter:

 **NOTE: This example shows three additional port groups added.**

```
symrpi -sid 124 -cluster CorpRPCluster environment -expand -nop
```

```
An RP 'Environment Expand operation is
in progress for cluster 'CorpRPCluster'. Please wait...
```

```
Setup and Expand RecoverPoint Environment.....Started.
Setup and Expand RecoverPoint Environment.....Done.
Setup and Expand RecoverPoint Environment.....Started.
Setup and Expand RecoverPoint Environment.....Done.
Setup and Expand RecoverPoint Environment.....Started.
Setup and Expand RecoverPoint Environment.....Done.
```

```
The RP 'Environment Expand' operation successfully executed for
cluster 'CorpRPCluster'.
```

RecoverPoint journal device deletion - description and actions

Description


The delete journal devices operation removes visibility of journal devices from the RP Appliance and deletes the devices.

 **NOTE: This operation can be run by the storage administrator on the Solutions Enabler control host or by the RP system.**

Actions

The delete journal operation performs the following actions:

- Validates that the specified RP Cluster exists on the target array, otherwise the operation fails.
- If removing specific Journal devices from RP storage group:
 - Validates that the device exists, otherwise the operation fails.
 - Validates that the device is a RP Journal Device, otherwise the operation fails.
 - Removes the device from the RP storage group that contains the journal devices.
 - Removes the devices from the RecoverPoint Journal storage group associated with the specified cluster.

 **NOTE: Even if the delete operation results in the removal of all Journal devices, the storage group is not be deleted.**

- Deletes the device(s).

Delete RecoverPoint journal devices

Description

The `symrpi delete -journal` command removes visibility of Journal devices from the RP Appliance and deletes devices. This operation requires the following security privileges:

- Access type – CREATEDV
- Authorization rights – Storage Admin

The command requires the following parameters:

- Target array (`-sid symID`)
- RP Cluster name (`-cluster ClusterName`)
- Devices to delete (`-dev <SymDevStart>:<SymDevEnd> | SymDevName`)

Syntax

To delete journal devices, use the following syntax:

```
symrpi -sid symID -cluster ClusterName delete -journal
```

Examples

To delete `-journal` devices 005B5 - 005B9 for cluster `CorpRPcluster` on array 124, enter:

```
symrpi -sid 124 -cluster CorpRPcluster delete -journal -dev 005B5:005B9
```

```
A RecoverPoint 'delete Journal' operation is  
in progress for cluster 'CorpRPcluster'. Please wait...
```

```
Analyze Configuration.....Validated.  
Delete Devices.....Started.  
Delete Devices.....Done.
```

```
The RecoverPoint 'Delete Journal' operation successfully executed  
for cluster 'CorpRPcluster'.
```

RecoverPoint device protection - description and actions

Description

The protect storage group operation protects the external storage group devices on the RP Cluster for use by the RP Appliance.

NOTE: This operation is issued by the storage admin on the Solutions Enabler control host to initially protect devices in a storage group. It can also be issued, either by the storage admin or by the RP Appliance, on a previously protected storage group to protect newly added devices.

Actions

The storage group protect operation performs the following actions:

- Validates that the cluster exists on the target array, otherwise the operation fails.
- Examines each device in the production storage group, determines if it is eligible for RP protection, then constructs a device list for RP tagging.
- Validates that there is enough room in at least one of the RP storage groups for the devices in the device list, otherwise the operation fails.
- Sets the `RP_External` tag on the device.

Rules and Restrictions

For devices to be eligible for RP protection the following restrictions apply. If any devices in the storage group do not meet these requirements the protect operation fails:

- Must have FBA emulation.
- Cannot have CKD, Celerra-FBA, or AS400 emulation
- Cannot be an Open Replicator control device
- Cannot be part of a SRDF STAR configuration
- Cannot be SRDF R1, R11, R2, R21, or R22 devices
- Cannot be part of a SRDF/Metro configuration
- Cannot be part of a NDM migration (neither source or target)
- Cannot have BCV attribute
- Cannot be encapsulated
- Cannot be a Data Domain device

NOTE: Any devices that are dedicated gatekeeper devices (< 20 cylinders) are not eligible for RP protection.

The protect operation will also fail for the following reasons:

- Specified storage group does not exist
- Specified storage group is empty
- Specified storage group is a child storage group in a cascaded relationship
- Specified storage group contains devices tagged *RP_Internal*
- Specified storage group contains RP protected devices that belong to another RP protected RP Cluster
- Specified storage group has more than 1024 devices, however this limit check is ignored if the storage group protect command is issued from the RP software

Protect devices for RecoverPoint

Description

The `symrpi protect` command protects the external storage group devices to the RP Cluster for use by the RP Appliance. This operation requires the following security privileges:

- Access type – RPA (for all devices in the storage group)
- Authorization rights – Storage Admin

The command requires the following parameters:

- Target array (`-sid symID`)
- RP Cluster name (`-cluster ClusterName`)
- Specified storage group for protection (`-sg SgName`)

Syntax

To RP protect devices, use the following syntax:

```
symrpi -sid symID -cluster ClusterName protect -sg SgName
```

Examples

To protect devices in storage group `MyCorp_Production` for `CorpRPCluster` on array 124, enter:

```
symrpi -sid 124 -cluster CorpRPCluster protect -sg MyCorp_Production -nop
```

```
A RecoverPoint 'SG Protect' operation is  
in progress for cluster 'CorpRPCluster'. Please wait...
```

```
Analyze Configuration.....Validated.  
Protect Storage Group.....Started.  
Protect Storage Group.....Done.
```

```
The RecoverPoint 'SG Protect' operation successfully executed for  
Cluster 'CorpRPCluster'.
```


RecoverPoint device protection removal - description and actions

Description

The unprotect storage group operation removes protection for the devices in a specified storage group on the RP Cluster.

NOTE: This operation is restricted and must be run by the storage administrator on the Solutions Enabler control host. It cannot be executed by the RP system.

This action can be used with a Symforce flag to remove storage groups from RP protection, when the RP Appliance is not available.

Actions

The storage group unprotect operation performs the following actions:

- Validates that the cluster exists on the target array, otherwise the operation fails.
- Validates that the specified storage group exists on the target array, otherwise the operation fails.
- Validates that the device is tagged as *RP_External*, otherwise the operation fails if device is tagged as *RP_Internal* or has no tagging.
- Removes the *RP_External* from the device.
- Removes the devices from the RP storage group that contains the production device.

Unprotect devices for RecoverPoint

Description

The `symrpi unprotect` command removes protection, for the devices in a specified storage group on the RP Cluster. This operation requires the following security privileges:

- Access type – RPA (for all devices in the storage group)
- Authorization rights – Storage Admin

The command requires the following parameters:

- Target array (`-sid symID`)
- RP Cluster name (`-cluster ClusterName`)
- Specified storage group or devices to unprotect (`-sg SgName`) or (`-devs SymDevStart:SymDevEnd |SymDevName...`)

Syntax

To unprotect devices for RP, use the following syntax:

```
symrpi -sid symID -cluster ClusterName unprotect -sg SgName
```

Options

-symforce

Required with this command to force removal of all RP resources associated with the storage group when the RP Appliance is not available.

Examples

To remove protection for devices in storage group `MyCorp_Production` for `CorpRPCluster` on array 124, enter:

```
symrpi -sid 124 -cluster CorpRPCluster unprotect -sg MyCorp_Production -nop
```

```
A RecoverPoint 'SG Unprotect' operation is  
in progress for cluster 'CorpRPCluster'. Please wait...
```

```
Analyze Configuration.....Validated.  
Unprotect Storage Group.....Started.  
Unprotect Storage Group.....Done.
```

```
The RecoverPoint 'SG Unprotect' operation successfully executed for  
Cluster 'CorpRPCluster'.
```

RecoverPoint repository creation - description and actions

Description

The create repository operation creates a 5.7 GB Repository device, tags it as an internal device, and adds it to the RP storage group.

NOTE: This operation can be run by the storage administrator on the Solutions Enabler control host or by the RP system.

Actions

The create repository operation performs the following actions:

- Validates that the cluster exists on the target array, otherwise the operation fails.
- Validates that there is enough space in at least one of the RecoverPoint storage groups to hold the device. Saves the internal identifier that the array assigned to that RecoverPoint storage group, and fails if the storage group does not exist.
- Validates that there is adequate space in at least one of the collections of Recover Point Storage Groups, otherwise the operation fails.
- Creates a 5.7 Gb repository file tags as *RP_Internal* and adds it to the RP storage group with internal identifier.

Create repository for RecoverPoint

Description

The `symrpi create -repository` command adds a 5.7 GB repository device to the RecoverPoint Cluster for use by the RecoverPoint system, and is valid even if multiple repository devices already exist.

This operation requires the following security privileges:

- Access type – CREATEDV
- Authorization rights – Storage Admin

The command requires the following parameters:

- Target array (`-sid symID`)
- RP Cluster name (`-cluster ClusterName`)

Syntax

To recreate a repository for RP, use the following syntax:

```
symrpi -sid symID -cluster ClusterName create -repository -nop
```

Examples

To create a Repository device for cluster `CorpRPcluster` on array 124, enter:

```
symrpi -sid 124 -cluster CorpRPcluster create -repository -nop
```

```
A RecoverPoint 'Repository Create' operation is  
in progress for cluster 'CorpRPcluster'. Please wait...
```

```
Analyze Configuration.....Validated.  
Create Repository Device RecoverPoint.....Started.  
Create Repository Device RecoverPoint.....Done.
```

```
Devices Created: 005B6
```

```
The RecoverPoint 'Repository Create' operation successfully executed for  
cluster 'CorpRPcluster'.
```

RecoverPoint environment removal - description and actions

Description

The environment remove operation removes the RP environment from the array.

NOTE: This operation is restricted and must be run by the storage administrator on the Solutions Enabler control host. It cannot be executed by the RP system.

Prerequisite cleanup actions

The following actions from the RP System are required prior to the environment remove operation, otherwise the operation fails:

- Verify that there is no Repository volume associated with the RP Cluster on the array.
- Remove all journal volumes associated with the RP Cluster on the array.
- Remove all replication infrastructure (SDDF and SnapVX sessions) associated with devices that were protected by the RP Cluster on the array.
- Remove RP protection from all production devices that were protected by the RP Cluster on the array.
- Remove all Access TDEVs on the array that are associated with the RP Cluster. Creating Access TDEVs is an operation restricted to only the RP system and cannot be executed by Solutions Enabler. These devices are used by RP during the replication process.

Actions

The environment remove operation performs the following actions:

NOTE: When the RP Appliance is not available, the environment remove operation can be used with a Symforce flag, which bypasses the validation checks. Proceed with caution when using this option as data loss is possible.

- Validates that the cluster exists on the target array.
- Cleans up RP storage groups for all devices that have *RP_External* tag – removes replication sessions, removes device from the RP storage group, removes RP tag. Requires the *-symforce* flag.
- Cleans up RP storage groups for all devices that have *RP_Internal* tag and are an Access TDEV or Journal device – removes device from the RP storage group and deletes device. Requires the *-symforce* flag.
- Removes Repository devices from RP storage groups and deletes the devices. Requires the *-symforce* flag.
- Removes all masking views, initiator groups, and storage groups associated with the the RP cluster, retains initiator group *RP_<ClusterName>_IG*.
- Removes gatekeeper devices from RP storage groups and deletes gatekeeper devices.

Remove environment for RecoverPoint

Description

The `symrpi environment remove` command removes the RP environment from the array. This command is performed after all application storage groups have been unprotected from the RP cluster.

NOTE: This operation can take a long time to execute depending on the number and size of the devices to be removed. A warning message displays when this remove action is requested.

This operation requires the following security privileges:

- Access type – CFGSYM
- Authorization rights – Storage Admin

The command requires the following parameters:

- Target array (*-sid symID*)
- RP Cluster name (*-cluster ClusterName*)

Syntax

To remove the RP environment on a target array, use the following syntax:

```
symrpi -sid symID -cluster ClusterName environment -remove -symforce -nop
```

Options

-symforce

Required with this command to force removal of all RP resources associated with the storage group when the RP Appliance is not available.

-nop

Requests no prompts are returned after the command is executed. The default is to prompt for command confirmation.

Examples

To remove the RP environment for cluster `CorpRPcluster` on array 124, enter:

```
symrpi -sid 124 -cluster CorpRPcluster environment -remove -symforce -nop
```

```
Warning: The 'Environment Remove' operation on cluster 'CorpRPcluster' may take a long time.
```

```
A RP 'Environment Remove' operation is  
in progress for cluster 'CorpRPcluster'. Please wait...
```

```
Analyze Configuration.....Validated.  
Remove Masking View(s).....Validated.  
Remove Storage Group(s).....Validated.  
Remove RecoverPoint Environment.....Started.  
Remove Masking View(s).....Started.  
Remove Masking View(s).....Done.  
Remove Storage Group(s).....Started.  
Remove Storage Group(s).....Done.  
Remove RecoverPoint Environment.....Done.
```

```
The RP 'Environment Remove' operation successfully executed for cluster 'CorpRPcluster'.
```

RecoverPoint integration reporting

The `symrpi list` command along with command options lists the following information:

- Summary information about each RP Cluster
- Specific cluster information – summary or details
- Specific production storage group information protected by a specific RP Cluster – summary or details

List RecoverPoint clusters

Description

The `symrpi list` command lists summary information (number of protected storage groups and total number of protected devices and journals) about all RP clusters configured on the array.

Syntax

To list RP clusters on target array, use the following syntax:

```
symrpi list -sid symID
```

Examples

To list the clusters on array 084, enter:

```
symrpi list -sid 084
```

Sample output

```
Symmetrix ID: 000197100084

                Protect
Cluster Name      SGs  Dev   Dev   Jrnl  Flags
                -----
Production_Cluster  3   182   182    2    .
Test_Cluster       37  1243  1197   112   R

Legend:
Flags:
(R)epository, . - N/A
```

List specific RecoverPoint cluster

Description

The `symrpi list -cluster` command lists summary information about a specific RP cluster, as well as summary information on each production storage group containing devices protected by the cluster. This command displays for each storage group count of devices in the group (sum of all devices in its children for a cascaded group), the count of devices protected under that storage group, and the number of data copies being maintained by RP for the group. The copy count is reported as `Mixed` if the count varies between devices within the group.

Syntax

To list information for an RP cluster, use the following syntax:

```
symrpi list -sid symID -cluster ClusterName
```

Options

-details

Provides additional information when one or more storage groups have the *Devices Removed* flag set. This option displays information on devices that have been removed from the storage group where they were RP protected.

Examples

To list information for `Production_Cluster` on array 084, enter:

```
symrpi list -sid 084 -cluster Production_Cluster
```

Sample output

```
Symmetrix ID      : 000198700084
RP Cluster        : Production_Cluster
Repository        : Yes
Journal Devices   : 2
Access TDevs     : 5
Repository Devices : 2
Protected Devices : 3

                Protect
Protected SG     Dev   Dev   Data  Flags
                -----
MyCorp_Production  35   35    2    ..
MyCorp_oracle_archive 100  100  Mixed ..
MyCorp_accts_payable  47   47    4    ..
```

```

-----
Device
-----
Name                Type
-----
05D8D              JOURNAL
05D8E              JOURNAL
05D8F              ATDEV
05D9A              ATDEV
05D9B              ATDEV
05D9C              ATDEV
05D9D              ATDEV
05DA2              REPOSITORY
05DA3              REPOSITORY
00100              PROTECTED
05D85              PROTECTED
05D86              PROTECTED

Legend:
Flags:
(D)evices: A - Added, R - Removed, . - N/A
(P)rotect: N - Needs Re-protect, . - N/A

```

Output with -detail option, which shows a set of devices (113A - 113F) that were removed from the storage group MyCorp_oracle_archive after it was RP-protected.

```

Symmetrix ID       : 000198700084
RP Cluster         : Production_Cluster
Repository         : Yes
Journal Devices   : 2
Access TDevs      : 5
Repository Devices : 2
Protected Devices : 3

                Protect
                Dev   Dev   Data  Flags
Protected SG   Count Count Copies DP
-----
MyCorp_Production      35    35     2  ..
MyCorp_oracle_archive  94   100     1  R.
MyCorp_accts_payable  47    47     4  ..

Misplaced
Devices Reason Protected SG
-----
  0113A No_SG MyCorp_oracle_archive
  0113B No_SG MyCorp_oracle_archive
  0113C No_SG MyCorp_oracle_archive
  0113D No_SG MyCorp_oracle_archive
  0113E No_SG MyCorp_oracle_archive
  0113F No_SG MyCorp_oracle_archive

-----
Device
-----
Name                Type
-----
05D8D              JOURNAL
05D8E              JOURNAL
05D8F              ATDEV
05D9A              ATDEV
05D9B              ATDEV
05D9C              ATDEV
05D9D              ATDEV
05DA2              REPOSITORY
05DA3              REPOSITORY
00100              PROTECTED
05D85              PROTECTED
05D86              PROTECTED

Legend:
Flags:

```

(D)evices: A - Added, R - Removed, - N/A
(P)rotect: N - Needs Re-protect, . - N/A

List Recover Point protected storage group

Description

The `symrpi list -cluster -sg` command lists summary information about a specific production storage group protected by a specific RP Cluster. This command displays summary information about the storage group and its protection properties, the number of devices within the group, the count of devices that are protected under the group, and for each device the number of data copies maintained by RP.

Syntax

To list information for an RP protected storage group, use the following syntax:

```
symrpi list -sid symID -cluster ClusterName -sg SGName -dev_info <journal | atdev | protected  
| repository | all>
```

Options

-detail

Provides the additional detail on each device within the storage group. This option displays the number of copies RecoverPoint maintains on the device and the current RecoverPoint status of the device, related to this storage group.

Examples

To list information for storage group `MyCorp_Production` `Production_Cluster` on array 084, enter:

```
symrpi list -sid 084 -cluster Production_Cluster -sg MyCorp_Production
```

Sample output

```
Symmetrix ID      : 000198700084  
RP Cluster       : Production_Cluster  
Storage Group    : MyCorp_Production  
Device Count     : 38  
Protected Device Count : 35  
Copies          : 2  
Flags           : Devices Added, Needs Re-Protect
```

Output with `-detail` option, which shows a set of devices added to the storage group `MyCorp_Production` that are not RP protected.

```
Symmetrix ID      : 000198700084  
RP Cluster       : Production_Cluster  
Storage Group    : MyCorp_Production  
Device Count     : 38  
Protected Device Count : 35  
Copies          : 2  
Flags           : Devices Added, Needs Re-Protect
```

Device List:

```
Device Copies Status
```

```
-----
```

```
....  
003BF          2 Protected  
003C0          2 Protected  
003EF          0 Unprotected  
003F0          0 Unprotected  
003F1          0 Unprotected  
00401          2 Protected
```

```
00402      2 Protected
.....
```

Output with `-dev_info all` option, which shows all of the devices added to the storage group on array 406.

```
Symmetrix ID : 000197800406
RP Cluster   : Production_Cluster
Repository   : Yes
Journal Devices : 5
Access TDevs : 5
Repository Devices : 2
Protected Devices : 3
Protect

Protected SG   Dev Count   Dev Count   Data Copies   Flags
-----
dev-50         3           3            0             ..

-----
Device
-----
Name          Type
-----
05D8A         JOURNAL
05D8B         JOURNAL
05D8C         JOURNAL
05D8D         JOURNAL
05D8E         JOURNAL
05D8F         ATDEV
05D9A         ATDEV
05D9B         ATDEV
05D9C         ATDEV
05D9D         ATDEV
05DA2         REPOSITORY
05DA3         REPOSITORY
00100         PROTECTED
05D85         PROTECTED
05D86         PROTECTED

Legend:
Flags:
(D)evices: A - Added, R - Removed, . - N/A
(P)rotect: N - Needs Re-protect, . - N/A
```

Report RecoverPoint device types

Description

The `symdev list -rp` command lists all the RP external or internal devices on a target array.

For an RP tagged device, the `symdev show` command lists the device tag type (external or internal).

Syntax

To list all RP devices (external or internal), use the following syntax:

```
symdev -sid SymID list -rp [external|internal]
```

Examples

To list all RP external devices on array 476, enter:

```
symdev -sid 476 list -rp external
```

To list the RP device tag type for RP device 1E3, enter:

```
symdev -sid 476 show 1E3
```


Sample output

For all RP external devices on array:

Symmetrix ID: 000196801476

Device Name		Dir		Device		
Sym	Physical	SA :P	Config	Attribute	Sts	Cap (MB)
00080	Not Visible	***:***	TDEV	N/Grp'd	RW	188
00081	Not Visible	***:***	TDEV	N/Grp'd	RW	188
00082	Not Visible	***:***	TDEV	N/Grp'd	RW	188
00083	Not Visible	***:***	TDEV	N/Grp'd	RW	188
00084	Not Visible	***:***	TDEV	N/Grp'd	RW	188
00085	Not Visible	***:***	TDEV	N/Grp'd	RW	188
00086	Not Visible	***:***	TDEV	N/Grp'd	RW	188
00087	Not Visible	***:***	TDEV	N/Grp'd	RW	188
00088	Not Visible	***:***	TDEV	N/Grp'd	RW	188
00089	Not Visible	***:***	TDEV	N/Grp'd	RW	188
000E8	Not Visible	***:***	TDEV	N/Grp'd	RW	188
001E3	Not Visible	***:***	TDEV	N/Grp'd	RW	3750
001E4	Not Visible	***:***	TDEV	N/Grp'd	RW	3750
001E5	Not Visible	***:***	TDEV	N/Grp'd	RW	3750
.....						

For RP device tag type for specific RP device:

```
Device Physical Name      : Not Visible

Device Symmetrix Name     : 001E3
Device Serial ID         : N/A
Symmetrix ID              : 000196801476

Number of RAID Groups    : 0
Encapsulated Device      : No
Encapsulated WWN         : N/A
Encapsulated Device Flags: None
. . .
Device Service State     : Normal

Device Status             : Ready           (RW)
Device SA Status         : Ready           (RW)
Device User Pinned       : False
Host Access Mode         : Active
Device Tag(s)            : RecoverPoint External

Extent Based Clone       : None
Snapvx Source            : False
Snapvx Target            : False
Data Destaged            : True
DIF1 Flag                 : False
Gatekeeper Device        : False
AS400 GK                 : False
Host Cache Registered    : False
Optimized Read Miss      : N/A
```

This chapter describes FAST.X concepts and explains how to configure and manage FAST.X using the SYMCLI.

Topics:

- FAST.X
- eDisks
- Create external disk group
- Delete an external disk group
- Add an eDisk
- Show external disk information
- Remove an eDisk
- List external spindle state
- Verify eDisk state
- Start drain on eDisk
- Stop drain on eDisk
- Activate eDisk

FAST.X

FAST.X attaches external storage to VMAX3 arrays (100K, 200K, and 400K) and 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 EMC storage array management. FAST.X uses DX director emulation. The DX director is transparent to other director emulations and operating array infrastructure, and allows the array to act on the external logical units as if they were internal physical drives.

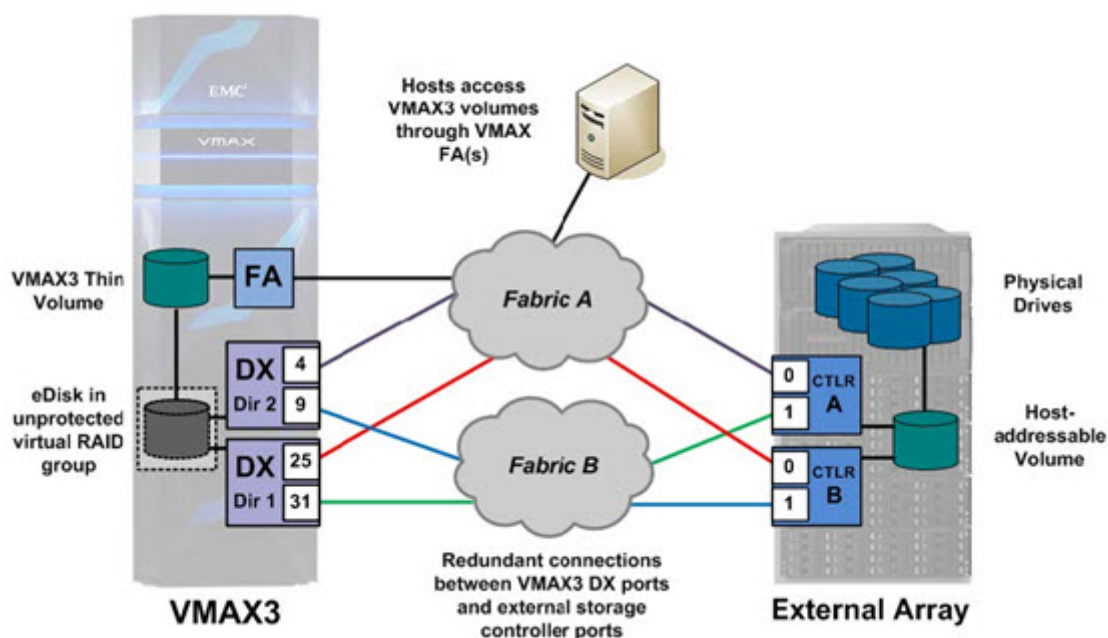


Figure 9. High-level overview of FAST.X environment

FAST.X overview

FAST.X external provisioning requires a FAST.X license and Advanced Suite license pack.

FAST.X for VMAX 100K, 200K, and 400K arrays, supports external provisioning for the following platforms:

- Data Domain eDisk encapsulation for EMC ProtectPoint – ProtectPoint allows for direct backup from VMAX3 arrays to Data Domain platforms. Support is provided for one external disk group and one associated pool to store external devices. The disk group and pool are created automatically during the VMAX array pre-configuration process and all encapsulated virtually-provisioned devices are added during this process.
- XtremIO, VNX, VMAX, Cloud Array and some third party arrays — Integration with Service Levels allows FAST management in these external arrays.

For more information on FAST.X refer to *EMC VMAX3 Family with HYPERMAX OS Product Guide*.

NOTE: VMAX 250F, 450F, 850F, and 950F arrays, do not support FAST.X, but do support external provisioning for Cloud Array and ProtectPoint.

Solutions Enabler features supported with FAST.X

Feature	Description
FAST VP	FAST management integrated with Service Levels.
VLUN migration	
Open Replicator	<ul style="list-style-type: none"> • Full support for externally provisioned device for ORS, and RP operations. • Support for encapsulated devices that are not geometry limited for only ORS and RP operations. • Encapsulated devices that are geometry limited are not supported for ORS pull operations.
SRDF	<ul style="list-style-type: none"> • All SRDF operations are supported for externally provisioned devices with any RDF personality (R1, R2, etc.). • All SRDF operations are supported for encapsulated devices that are not geometry limited with any RDF personality, (R1, R2, etc.). • Encapsulated devices that are geometry limited are supported for SRDF migration operations only. Support is limited to R1 devices. The rules of operation (R2 larger than R1) apply. • Encapsulated Data Domain devices used for ProtectPoint cannot be part of any SRDF device pair.
SRDF/SAR	<p>The symreplicate CLI supports both encapsulated and externally provisioned devices, as follows:</p> <ul style="list-style-type: none"> • Supports externally provisioned devices. • Supports encapsulated devices that are not geometry limited. • Does not support encapsulated devices that are geometry limited for any SAR operation. <p>The symrecover CLI supports both encapsulated and externally provisioned devices, as follows:</p> <ul style="list-style-type: none"> • Supports externally provisioned devices. • Supports encapsulated devices that are not geometry limited. • Does not support encapsulated devices that are geometry limited. • Encapsulated Data Domain devices used for ProtectPoint cannot be part of any SRDF device pair.
SRDF/Star	<ul style="list-style-type: none"> • Supports externally provisioned devices. • Supports encapsulated devices that are not geometry limited. • Does not support encapsulated devices that are geometry limited for any Star operation. • Encapsulated Data Domain devices used for ProtectPoint cannot be part of any SRDF device pair.

Feature	Description
TimeFinder	<ul style="list-style-type: none"> • Supports externally provisioned devices as source and target devices for TimeFinder/Snap, Clone, and Mirror (TF/Clone Emulation) operations. • Supports encapsulated devices that are not geometry limited as source and target devices for TimeFinder/Clone, and TimeFinder/Mirror (TF/Clone Emulation) operations. • Supports encapsulated devices that are geometry limited only for TimeFinder/Clone operations. These devices can be used as a source devices and the rules for operations to larger target devices apply. • Encapsulated Data Domain devices used for ProtectPoint cannot be part of Clone/Mirror sessions.

eDisks

With a VMAX3 array attached to external storage, FAST.X virtualizes an external array's SCSI logical units as VMAX disks called eDisks. eDisks have two modes of operation:

- Encapsulation — Preserves existing data on external arrays and accesses it through VMAX devices. These devices are called *encapsulated devices*.
- External Provisioning — Uses external storage as raw capacity for new VMAX devices. These devices are called *externally provisioned devices*. Existing data on the external devices is deleted when they are externally provisioned.

Restrictions

The following restrictions apply to eDisks:

- Must be unprotected devices. The RAID protection scheme of eDisks is dependent on the external array.
- Cannot be AS400, CKD, or gatekeeper devices.
- Cannot be used as VAULT, SFS, or ACLX devices.

eDisk encapsulation

eDisk encapsulation has two modes of operation:

- Encapsulation for disk group provisioning (DP encapsulation) — The eDisk is encapsulated and exported from the VMAX array as disk group provisioned devices.
- Encapsulation for virtual provisioning (VP encapsulation) — The eDisk is encapsulated and exported from the VMAX array as thin devices.

For both modes, the array automatically creates the necessary VMAX devices. If the eDisk is larger than the maximum VMAX device size or the configured minimum auto meta size, the array creates multiple VMAX devices to account for the full size of the eDisk. These VMAX devices are concatenated into a single concatenated meta device and allow access to the complete volume of data available from the eDisk.

Create external disk group

Description

eDisks require an external disk group (where eDisks are added).

Syntax

To create an external disk group, use the following syntax:

```
symconfigure create disk_group DskGrpName disk_location = <external>
```

 **NOTE:** The disk group name serves as an additional identification mechanism and does not replace disk group numbers.

Example

To create a disk group for a storage pool named `hr_sg_pool`, enter:

```
symconfigure -sid 2012 -cmd "create disk_group hr_disk_group disk_location = external;" commit
```

Restrictions

The following restrictions apply when creating disk groups.

- Creating external disk groups is not supported in HYPERMAX OS 5977 and later.
- The maximum number of external disk groups allowed on an array is 512.
- A disk group cannot be created and deleted in the same session.
- The specified disk group name must be unique; the check for uniqueness is case insensitive.
- The maximum characters allowed in a disk group name is 32. Only alphanumeric characters, hyphens "-" and underscores "_" are allowed. The name cannot start with a hyphen "-" or underscore "_".
- The disk group name given cannot have the same format as the default disk group names, for example DISK_GROUP_001.

Delete an external disk group

Syntax

To delete a disk group by disk group number or by disk group name, use the following syntax:

```
symconfigure delete disk_group <DskGrpNum | name:DskGrpName>;
```

Examples

To delete a disk group with the name `hr_disk_group`, enter:

```
symconfigure -sid 2012 -cmd "delete disk_group name:hr_disk_group;" commit
```

Restrictions

The following restrictions apply when deleting disk groups:

- The disk group must be empty.
- The disk group name or number must exist.
- Only external disk groups can be deleted.

The following rules apply to physical disks and disk groups when using this feature:

- An external disk group can only contain external disks.
- Every external disk in the array must belong to an external disk group.

Add an eDisk

Description

When adding an eDisk to a VMAX array, it must be added to an existing external disk group.

NOTE: Use the `sysmsan` command to obtain the WWN of the external LUN to specify in the `add external_disk` command. Refer to the *EMC Solutions Enabler SYMCLI Command Reference* for the `sysmsan` manpage.

Syntax

For HYPERMAX OS 5977 lower than Q12016SR and for ProtectPoint on Data Domain platforms, to add an eDisk use the following syntax:

```
add external_disk wwn=<wwn> encapsulate_data=<YES | NO <SRP=<SRPName>>> [dir=<Director_num>];
```

For HYPERMAX OS 5977 Q12016SR and higher to add an eDisk, use the following syntax:

NOTE: Other than ProtectPoint, this syntax with the `keep_data` option, is used when working with all supported external storage platforms.

```
add external_disk wwn=<wwn> encapsulate_data=<YES | NO  
<keep_data=<YES [sg=<sgname>] | NO [SRP=<SRPName>]>>> [dir=<Director_num>];
```

Options

`encapsulate_data`

Set this option to `YES` to encapsulate data on external LUNs. For HYPERMAX OS 5977 Q12016SR or higher, when adding an external disk `encapsulate_data` is used ONLY when configuring eDisks on Data Domain for use with ProtectPoint.

`keep_data` (HYPERMAX OS 5977 Q12016SR and higher)

This option can only be specified when `encapsulate_data` is set to `NO`. Set this option to `YES` to retain existing data on external LUNs. If set to `YES`, accepts a Storage Group name. If set to `NO`, accepts a valid SRP configured on the array. If no SRP is specified then `DEFAULT_SRP` is used.

Examples

For HYPERMAX OS 5977, to virtualize external LUN with WWN 600219800002DDA8B into SRP "EXTERNAL_SRP", enter:

```
symconfigure -sid 087 commit -cmd "add external_disk wwn=600219800002DDA8B  
encapsulate_data=NO SRP=EXTERNAL_SRP;"
```

For HYPERMAX OS 5977 Q12016SR and higher, to virtualize external LUN with WWN 600219800002DDA8B and retain existing data (incorporation), enter:

```
symconfigure -sid 087 commit -cmd "add external_disk wwn=600219800002DDA8B  
encapsulate_data=NO keep_data=YES sg=SG1;"
```

Rules and restrictions for adding eDisks

HYPERMAX OS 5977 or higher:

- Required Access type: CFGSYM.
- Required Authorization Rights: Storage Admin.
- External LUN size must be between 42 cylinders and 16TB.
- Requires the Advanced Suite license pack and FAST.X license to add eDisks.

HYPERMAX 5977Q32015SR or higher

All external LUNs from the same array, virtualized in the same SRP, must be the same size.

HYPERMAX OS 5977 Q32015SR or lower

A total of 2048 eDisks per system can be virtualized on the array. This includes both encapsulation and external provisioning.

HYPERMAX 5977 Q12016SR or higher:

- If eDisk is added with `keep_data` set to `YES`, data can remain on external array indefinitely. To remove external LUN, first use the `start drain` on `external_disk` command to drain the data to VMAX. When drain is complete, remove the external LUN.
- If `keep_data` is set to `YES`:
 - Specified Storage Group cannot contain CKD devices.
 - Specified Storage Group, if empty, must have SRP configured as FBA DEFAULT on the array.
 - Specified Storage Group, if empty, cannot be set to SRP configured as CKD DEFAULT.
- A total of 2048 eDisks per engine can be virtualized on the array. This includes both encapsulation and external provisioning.

Show external disk information

Examples

To show external disk information for spindle ID 1E05, enter:

```
symdisk show -spid 1E05 -sid 432
```

Sample output

```
Symmetrix ID           : 000194900432
Director               : DX-2F
Interface              : N/A
Target ID              : N/A
Spindle ID             : 1E05
External WWN           : 60000970000184700306533030314345
External Array ID     : 000194900345
External Device Name   : 00040
Disk Group Number      : 512
Disk Group Name        : DISK_GROUP_512
Disk Location          : External
...

Hypers (1):

{
#   Vol   Emulation      Dev   Type           Mir  Mbr  Sts  Cap (MB)
-----
  1 N/A   FBA                00510 Ext-Data        1   1   RW   17261 }
}
```

Remove an eDisk

Description

For externally provisioning eDisks, the VMAX devices on the eDisk must be drained before it can be removed. See [Start drain on eDisk](#) on page 273 on how to start draining an external disk.

For encapsulated eDisks, the devices on the eDisk must be unmapped and not be part of a migration, local replication, remote replication, or ORS session. RAID groups and any disk group provisioned devices or DATA devices created during the encapsulation are removed. Thin devices are unbound but they are not removed.

Syntax

To remove an eDisk, use the following syntax:

```
remove external_disk <wwn=<wwn> | spid=<SpindleID>> [force_remove=<YES | NO>];
```

Options

force_remove

This option forces the removal of an external encapsulated disk when data is still in the VMAX cache. This can occur when the device has iVTOC and write pending tracks.

Examples

To remove an eDisk specifying the eDisks WWN, enter:

```
symconfigure commit -cmd "remove external_disk wwn=60000970000184700306533030314345;"
```

To remove an eDisk using the specifying the spindle ID, enter:

```
symconfigure commit -cmd "remove external_disk spid=2256;"
```

Rules and restrictions for removing eDisk

- Required Access type: CFGSYM
- Required Authorization Rights: Storage Admin
- HYPERMAX OS 5977 requires Advanced Suite license pack and FAST.X license.
- The specified WWN or spindle ID must represent an external disk.
- The specified external disk must exist.
- The specified WWN or spindle ID must represent an external disk virtualized for external provisioning.
- The specified external disk must be in the *Drained* state.

List external spindle state

Syntax

To list the state of an external spindle, use the following syntax:

```
symdisk list -external -spindle -state [-spid <SymID>]
```

To list external spindle states for an array, use the following syntax:

```
symdisk list -external -spindle -state [-sid <#>]
```

Examples

To list external spindle states for array 084, enter:

```
symdisk list -external -state -spindle -sid 084
```

Sample output

```
Symmetrix ID : 000197100084
```

Spindle	FLG		Product	Array ID	State	Drained (%)
	T	Vendor				
8000	P	EMC	XtremIO	000012345678	Active	N/A
8001	P	EMC	XtremIO	000012345678	Draining	50
8002	P	EMC	XtremIO	000012345678	Drained	100

Legend:

Flags:

(T)ype: E - Encapsulated, P - External Provisioning

Output descriptions for State values:

- **Active** — At least one data device associated with the eDisk is enabled.
- **Draining** — At least one data device associated with the eDisk is draining.
- **Drained** — All data devices associated with the eDisk are disabled and not draining.
- **Drained (%)** — If the eDisk is fully allocated then the Drained % starts from 0. Otherwise it starts from a value depending upon the allocated tracks on the external disk. For example if the external disk is 100% allocated then the Drained % value starts from approximately 0%. If it is 70% allocated then the Drained % value starts at approximately 30%.
- **Disabled** — If draining is stopped using `stop drain`, the eDisk state becomes disabled and is not available for allocations.

Verify eDisk state

Syntax

To verify the state of an external disk, use the following syntax:

 **NOTE: The ExternalWWN argument for the -spid option must refer to an eDisk.**

```
syndisk verify -external -sid <SymmID> [<-wwn=<ExternalWWN> | -spid=<SpindleID>>] [-i <Interval>] [-c <Count>] < -draining | -drained | -active | disabled >
```

Examples

To verify the *Active* state for spindle 8000, enter:

```
syndisk verify -external -sid 084 -spid 8000 -active
```

```
All disk(s) are in the 'Active' state.
```

To verify the *Draining* state for spindle 8000, enter:

```
syndisk verify -external -sid 084 -spid 8000 -draining
```

```
None of the disk(s) are in the "Draining" state.
```

Start drain on eDisk

Description

The `start drain` command starts the background draining process. Use the `syndisk verify` or `syndisk list -state` commands to monitor the draining progress. (See [Verify eDisk state](#) on page 273 or [List external spindle state](#) on page 272). When draining is complete the external disk state is "Drained". If draining is stopped using the `stop drain` command, the external disk state changes from "Disabled" and the disk is not available for allocations. See [Stop drain on eDisk](#) on page 274 for `stop drain` command.

Syntax

To start the drain on an external disk, use the following syntax:

```
start drain on external_disk <wwn=<wwn> | spid=<SpindleID>>;
```

Examples

To start drain on external disk specifying the WWN, enter:

```
symconfigure -sid 087 commit -cmd "start drain on external_disk wwn=600219800002DDA8B;"
```

```
Execute a symconfigure operation for symmetrix '000197100087' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
  Establishing a configuration change session.....Established.
  Processing symmetrix 000197100087
  {
start drain on external_disk wwn=600219800002DDA8B;
  }
  Performing Access checks.....Allowed.
  . . .
  Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

To start drain on external disk specifying the spindle ID, enter:

```
symconfigure -sid 087 commit -cmd "start drain on spid=23;"
```

```
Execute a symconfigure operation for symmetrix '000197100087' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
  Establishing a configuration change session.....Established.
  Processing symmetrix 000197100087
  {
start drain on external_disk spid=23;
  }
  Performing Access checks.....Allowed.
  . . .
  Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute the `start drain` command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- Requires Advanced Suite license pack and FAST.X license.
- The specified WWN or spindle ID must represent an external disk.
- The specified external disk must exist.
- The specified WWN or spindle ID must represent an external disk virtualized for external provisioning.
- The specified external disk cannot be in the *Draining* or *Drained* state.
- Sufficient free capacity must be available in the SRP, for all the allocated tracks produced from draining the external disk.

Stop drain on eDisk

Description

Stopping the drain process, moves the external disk state to *Disabled*.

Syntax

To stop the drain on an external disk, use the following syntax:

```
stop drain on external_disk <wwn=<wwn> | spid=<SpindleID>>;
```

Examples

To stop drain on external disk specifying the WWN, enter:

```
symconfigure -sid 087 commit -cmd "stop drain on external_disk wwn=600219800002DDA8B;"
```

```
Execute a symconfigure operation for symmetrix '000197100087' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
  Establishing a configuration change session.....Established.
  Processing symmetrix 000197100087
  {
stop drain on external_disk wwn=600219800002DDA8B;
  }
  Performing Access checks.....Allowed.
  . . .
  Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

To stop drain on external disk specifying the spindle ID, enter:

```
symconfigure -sid 087 commit -cmd "stop drain on external_disk spid=23"
```

```
Execute a symconfigure operation for symmetrix '000197100087' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
  Establishing a configuration change session.....Established.
  Processing symmetrix 000197100087
  {
stop drain on external_disk spid=23;
  }
  Performing Access checks.....Allowed.
  . . .
  Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute the `start drain` command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- Requires Advanced Suite license pack and FAST.X license.
- The specified WWN or spindle ID must represent an external disk.
- The specified external disk must exist.
- The specified WWN or spindle ID must represent an external disk virtualized for external provisioning.
- The specified external disk cannot be in the *Drained* state.

Activate eDisk

Description

Activating an edisk is allowed when the external disk is in *DISABLED*, *DRAINED* or *DRAINING* state, and this action moves the edisk to *ACTIVE* state. Activating the edisk which is in *DRAINING* state stops the drain operation before activating.

Syntax

To activate an external disk, use the following syntax:

```
activate external_disk <wwn=<wwn> | spid=<SpindleID>>;
```

Examples

To activate an external disk specifying the WWN, enter:

```
symconfigure -sid 087 commit -cmd "activate external_disk wwn=600219800002DDA8B;"
```

```
Execute a symconfigure operation for symmetrix '000197100087' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
  Establishing a configuration change session.....Established.
  Processing symmetrix 000197100087
  {
activate external_disk wwn=600219800002DDA8B;
  }
  Performing Access checks.....Allowed.
  . . .
  Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

To activate an external disk specifying the spindle ID, enter:

```
symconfigure -sid 087 commit -cmd "activate external_disk spid=23;"
```

```
Execute a symconfigure operation for symmetrix '000197100087' (y/[n]) ? y
A Configuration Change operation is in progress. Please wait...
  Establishing a configuration change session.....Established.
  Processing symmetrix 000197100087
  {
activate external_disk spid=23;
  }
  Performing Access checks.....Allowed.
  . . .
  Terminating the configuration change session.....Done.
The configuration change session has successfully completed.
```

Restrictions

- The following security privileges are required to execute the `start drain` command:
 - Required Access type: CFGSYM
 - Required Authorization Rights: Storage Admin
- Requires Advanced Suite license pack and FAST.X license.
- The specified WWN or spindle ID must represent an external disk.
- The specified external disk must exist.
- The specified WWN or spindle ID must represent an external disk virtualized for external provisioning.
- The specified external disk cannot be in the *Active* state.

Device masking with Auto-Provisioning Groups

This chapter explains how to confine host access to array devices using auto-provisioning groups and the SYMCLI `symaccess` command.

Topics:

- [Auto-provisioning groups](#)
- [Create an auto-provisioning session](#)
- [Manage masking views](#)
- [Manage storage groups](#)
- [Manage port groups](#)
- [Manage initiator groups](#)
- [Display auto-provisioning group information](#)

Auto-provisioning groups

Auto-provisioning groups creates groups of host initiators, front-end ports, and logical devices. These groups are associated to form a masking view, where controls are managed. This feature reduces the number of commands needed for masking devices, and allows for easy management of the masking view.

Storage provisioning with the `symaccess` command creates a group of devices, a group of director ports, a group of host initiators, and with one command, associates them into a *masking view*. Once a masking view exists, devices, ports, and initiators are easily added or removed from these groups. The `symaccess` command can list login history and be used to manage the port flags on an initiator.

Provisioning limits

Provisioning object	Configuration limitations
Devices	Maximum 16K per each director Maximum 4K per each storage group
Initiator group	Maximum 64 initiator addresses (or 64 child IG names) per group NOTE: Using multiple child initiator groups in a cascaded initiator group with multiple initiators, allows a masking view to exceed the limit of 64 initiators.
Storage group	Maximum 16K storage groups per array Maximum 64 child storage groups per each parent storage group Maximum 4K storage groups with host I/O limits defined
Port group	Maximum 16K port groups per array Maximum 32 ports in a port group
Masking view	Maximum 64K masking views per array
LUN addresses	Maximum 4K LUN addresses per director port

Create a masking view

The steps for creating a masking view are:

1. Create a storage group (one or more devices).
2. Create a port group (one or more director/port combinations).
3. Create an initiator group (one or more host WWNs or iSCSIs).
4. Create a masking view containing the storage group, port group, and initiator group.

The devices are automatically masked and mapped when a masking view is created.

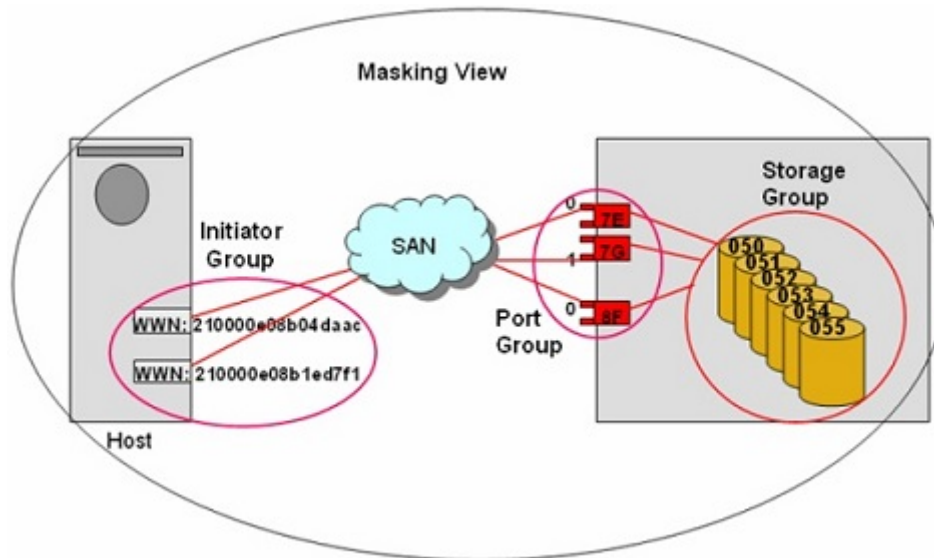


Figure 10. Masking view overview

After the masking view is created, any objects (devices, ports, initiators) added to a group automatically become part of the masking view.

Auto-provisioning session rollback

Auto-provisioning operations (create masking views, or add devices or ACLX-enabled ports to existing views) continue to map devices to ports, but if the masking view update fails these devices are unmapped when the session is rolled back.

Storage groups containing CKD devices must already be mapped, and must use the optional flag `-ckd`. Storage groups containing Celerra devices can be masked (and mapped) using the `-celerra` option. Storage groups containing devices tagged for RecoverPoint are masked and mapped using the `-rp` option.

HYPERMAX OS 5977 or higher

Explicit mapping and unmapping of ACLX-enabled ports using the `symconfigure` command is not supported. To make the ACLX device visible or to remove visibility from the ACLX device, modify the new `SHOW_ACLX_DEVICE` attribute on front end ports. By default, ACLX-enabled ports have the `SHOW_ACLX_DEVICE` attribute disabled, making the ACLX device not visible to hosts zoned to the port. One ACLX-enabled port is pre-configured with the `SHOW_ACLX_DEVICE` attribute enabled. During device remove, port remove, or masking view delete operations, the `-unmap` option is ignored.

Create an auto-provisioning session

This section details the following steps for creating an auto-provisioning session:

1. Discover the host HBAs.
2. List (identify) the host HBAs.
3. Create groups and views.

Discover host HBAs

Description

During the initial array setup, a search, from the controlling host, for each HBA connected to array devices is performed using the `symaccess discover` command. The `symaccess discover` command sends information about this connection back to its host system. The discover command is the primary mechanism by which hosts, other than the control station, identify their paths to the array.

Examples

To discover host HBAs: enter

```
symaccess discover hba
```

When the `symaccess discover` command finds a host HBA, it reads the login history table and performs the following:

1. Creates an ASCII alias and writes it to the login history table.
NOTE: There is a `-rename` option that is used with the `symaccess discover` command, to force a write of the discovered hostname/HBA name (or IP address) to the login history table and the initiator group.
2. Prints the initiator identifier (WWN/iSCSI) of the HBAs connected to the masked channel and the array.

List host HBAs

Description

The `symaccess discover` and `symaccess list` commands do not use the ACLX device to determine the available connections. Therefore, connections available through all gatekeeper devices are listed regardless of the host visibility of the ACLX device. For the `symaccess list` command, the connections displayed include the paths detected through either the ACLX device or any gatekeeper device. If multiple gatekeepers, including the ACLX device, are available for detecting connections, only one entry is displayed for each unique connection to the array.

NOTE: An iSCSI initiator cannot log in to the array until it belongs to a masking view that includes that specific port on the array.

Examples

To list host HBAs, enter:

```
symaccess list hba
```

Sample output

Shows only one unique connection through multiple gatekeepers.

Identifier	Physical Device Path	Symmetrix ID	Dir:Port
10000000c99527b2	/dev/sdh	000197100086	01E:009
	/dev/sdw	000197300005	01E:008
	/dev/sdb	000198700030	01E:000

Create groups and views

The `symaccess` command is used to create the initiator groups, port groups, and storage groups that make up the masking view.

Restrictions and limitations

The following are restrictions and limitations for the `symaccess create` command:

- Storage Admin rights are required.

- The command fails if the device list contains both FBA and CKD devices.
- The command fails if the list of child SGs contain both FBA and CKD devices.
- The command fails if the device list contains CKD devices and a workload is given.
- The command fails if the device list contains CKD devices and the Service Level or SRP given do not support CKD devices.

Create storage groups

Description

The `symaccess` and the `symsg` commands support storage group operations. Use the `symaccess` command for creating storage groups by specifying a range of devices, a list of devices, a device group, a storage group, or a device file.

NOTE: Host I/O Limits can be set for a storage group. Host I/O Limits are settings that limit the amount of front-end (FE) bandwidth (MBs) and I/Os per second (IOPs) that are consumed by a set of array devices over a set of director ports. Host I/O Limits for storage groups on page 107 provides feature details and restrictions.

Syntax

To create a storage group, use the following syntax:

```
symaccess -sid SymmID create -name GroupName -type storage
[devs SymDevStart:SymDevEnd |
  SymDevName[, SymDevName[, SymDevName...]] |
  <-g DgName [-std] [-bcv] [-vdev] [-tgt]> |
  <sg SgName[, SgName1, SgName2, ., SgName]>
  <-file DeviceFileName [src] [tgt]>
  [-reserve_id ResvID[, ResvID[, ResvID...]]]
```

NOTE: LUNs are assigned by the array when the masking view is created, not when devices are added to the storage group.

Examples

To create a storage group `SG_1` and add device range `050:055`, enter:

```
symaccess create -sid 458 -name SG_1 -type storage devs 050:055
```

Storage group and group name rules:

- Maximum group name length is 64 characters and are not case sensitive.
- Group Name must begin with an alphanumeric character, and embedded hyphens and underscore characters are valid.
- Group names must be unique per group type, but different group types can share the same name. For example, a storage group, a port group, and an initiator group can all have the name `Financial_DB`. However, two storage groups cannot be named `Financial_DB`.
- Device reservations are enforced whenever devices are added to a storage group.

Create port groups

Description

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

Syntax

To create a port group, use the following syntax:

```
symaccess -sid SymmID -name GroupName -type port
create [-dirport Dir:Port[, Dir:Port...]]
```


Examples

To create a port group PG_1 with three front-end ports, enter:

```
symaccess create -sid 458 -name PG_1 -type port -dirport 7E:0,7G:1,8F:0
```

Create initiator groups

Description

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. Initiator groups are created using the HBA's WWN, iSCSI, a file containing WWNs or iSCSI names, or another initiator group name.

NOTE: Different types of initiators (external Fibre Channel WWNs, internal guest Fibre Channel WWNs and iSCSI IQNs) cannot be mixed within a single Initiator Group. In addition, all child IG names added to a parent initiator group must contain the same Initiator type.

Syntax

To create an initiator group, use the following syntax:

```
symaccess -sid SymmID  
create -name GroupName -type initiator  
      [-consistent_lun]  
      [-wwn wwn | -iscsi iscsi |  
      -file InitiatorFilename | -ig InitiatorGroupName ]
```

Options

-consistent_lun

Use this option if the devices of a storage group (in a view) need to be seen on the same LUN on all ports of the port group. If the `-consistent_lun` option is set on the initiator group, the host LUN number assigned to devices is the same for the ports on the HBA. If this option is not set, the system will choose the first available LUN on each individual port.

Examples

To create an initiator group, IG_1, enter:

```
symaccess create -sid 458 -name IG_1 -type initiator -file IG_1
```

The file IG_1 contains:

```
wwn:210000e08b04daac
```

Create a masking view

Description

A masking view is a container of a storage group, a port group, and an initiator group, and makes the storage group visible to the host. Devices are masked and mapped automatically. The groups must contain some devices entries.




Figure 11. Masking view MV_1

Volume dynamic addressing is enabled by default. The array assigns the next available LUN address on the FA port when the masking view is created. The LUN assigned on the FA port will not necessarily match the masking LUN.

When you create a masking view, if Host I/O Limits have been set for the storage group, they become active with the `symaccess create view` command.

NOTE: Encapsulated devices that are being used as TimeFinder SnapVX source or target devices cannot be masked to a host.

Syntax

To create a masking view, use the following syntax:

```
symaccess create view -name <MaskingViewName> -sg <StorageGroupName> -pg <PortGroupName> -ig <InitiatorGroupName> [-lun <Addr>]
```

Options

`-lun <Addr>`

Specifies the LUN address for the devices being added to the masking view. When creating a new masking view using groups, only one LUN address is allowed. Multiple LUNs are allowed when using a device list, port list, and HBA list to create a view.

Examples

To create a masking view with storage group SG_1, port group PG_1, initiator group IG_1, and LUN 6, enter:

```
symaccess -sid 458 create view -name MV_1 -sg SG_1 -pg PG_1 -ig IG_1 -lun 6
```

After creating a masking view:

If additional storage is needed, devices added to the storage group are automatically masked and mapped in the masking view. This also applies to front-end ports and host initiators.

For managing other storage, create a second storage group and then create a masking view using the same port group and initiator group. The same number of LUNs can be supplied. Supplying this value is optional and the corresponding input flag should be supplied when it is given.

In a clustered environment, some devices may be seen by the entire cluster, but gatekeeper devices may only need to be seen by individual hosts.

Manage masking views

This section explains how to perform the following management operations on masking views:

- Delete masking views
- Name groups and views
- Back up and restore views
- Copy groups and views

Delete masking views

Syntax



NOTE:

When a masking view is deleted, all groups in the masking view remain intact. Any device reservations continue to be enforced when a masking view is deleted.

To delete a masking view, use the following syntax:

```
symaccess -sid SymmID delete view -name ViewName
```

Examples

For example, to delete masking view MV_1, on array 458, enter:

```
symaccess -sid 458 delete view -name mv_1
```

Name groups and masking views

Syntax

To rename a storage group, port group, or an initiator group, using the following syntax:

```
symaccess -sid SymmID rename -name GroupName  
-type <storage | port | initiator> -new_name NewGroupName
```

To rename a masking view, using the following syntax:

```
symaccess -sid SymmID rename view -name ViewName  
-new_name NewViewName
```

If the new name already exists, an error returns.

Storage group and masking view naming rules:

- Maximum group name length is 64 characters and are not case sensitive.
- Group Name must begin with an alphanumeric character, and embedded hyphens and underscore characters are valid.
- Group names must be unique per group type, but different group types can share the same name. For example, a storage group, a port group, and an initiator group can all have the name `Financial_DB`. However, two storage groups cannot be named `Financial_DB`.
- Renaming of an initiator group is propagated to the higher group if the group is cascaded.

Back up and restore masking views

Description

The masking views, including storage groups, port groups, and initiator groups can be backed up to and restored from a file.

Syntax

To backup the masking views to a file, use the following syntax:

```
symaccess -sid SymmID -f BackupFilename [-noprompt]
  backup
  restore [-remove_ckd] [-unused_sgs] [-disassociate]
```

The `symaccess` command validates the consistency of the Auto-provisioning data before the `backup` or `restore` actions are performed.

Options

- noprompt**
Eliminates the prompt for confirmation of the operation.
- remove_ckd**
Skips all CKD devices within the backup, allowing the backup to be restored if the CKD devices are no longer mapped.
- disassociate**
Disassociates the storage group from a FAST policy if the storage group contains invalid devices for FAST.

Examples

```
symaccess backup -sid 266 -f aclx_backup -nop
```

Example output with consistency errors:

```
Starting a backup operation.....
There are inconsistencies in the masking database. The operation cannot be performed.
```

Example output with no consistency errors:

```
Starting a backup operation.....
The masking data on Symmetrix 000192600266 was backed up to file aclx_backup.
```

Copy groups and masking views

Description

Groups and a copy of a storage, port, or initiator group, or a complete masking view can be copied from one array to another. When copying, any child view or cascaded initiator group are included in the copy action.

Syntax

To copy groups or views from one array to another, use the following syntax:

```
symaccess -sid SymmID -target_sid SymmID copy -name GroupName -type storage
  [-reserve_id ResvID [, ResvID [, ResvID ...]]]
copy -name GroupName -type initiator | port
copy -name ViewName view [-ckd] [-celerra] [-rp]
  [-reserve_id ResvID [, ResvID [, ResvID ...]]]
```

Options

- `-reserve_id`
Includes any device reservations.
- `-ckd`
Specifies that the view contains CKD devices
- `-celerra`
Specifies that the view contains Celerra devices (the devices will also be mapped).
- `-rp`
Includes devices that have been tagged for RecoverPoint.

Examples

For example, to copy masking view `mv_1` from array 207 to array 123, enter:

```
symaccess -sid 207 -target_sid 123 copy -name mv_1 view
```

Manage storage groups

After creating a storage group, as explained in [Create storage groups](#) on page 280, the following actions can be performed:

- Add devices
- Remove devices
- Rename a storage group
- Delete storage groups
- List and show storage group information

Add devices to storage group

Description

A storage group can contain up to 4k array device numbers, and devices can belong to more than one storage group. When adding devices specify the device names, a range of devices, a list of devices in a device group, or devices in a device file. Device reservations are still enforced when they are added to a storage group.

NOTE: Solutions Enabler supports adding storage groups (`symaccess ad sg`) that have Host I/O Limits set to a parent storage group, that is in a view using a port group with FCoE ports.

Syntax

To add devices to an existing storage group, use the following syntax:

```
symaccess -sid SymmID -name GroupName -type storage  
          [-reserve_id ResvID [, ResvID [, ResvID...]]]  
          [-ckd] [-celerra] [-rp]  
  
add devs SymDevStart:SymDevEnd [-lun Addr] |SymDevName [-lun Addr] |  
        SymDevName, SymDevName, SymDevName...  
        [-lun Addr | -lun Addr, Addr, Addr...]  
add -g DgName [-std] [-bcv] [-vdev] [-tgt] [-lun Addr]  
add -file DeviceFileName [src] [tgt] [-lun Addr]  
add sg SgName [, SgName1, SgName2, .., SgNameN]  
      [-lun Addr, Addr, Addr...]
```

Options

- `-ckd`
Adds CKD devices to a storage group.

-celerra

Adds (and maps) Celerra devices to a storage group.

-rp

Adds devices tagged for RecoverPoint to a storage group.

LUN address designation

When devices are added at the storage group creation time, do not specify a LUN address. The LUN address is determined when the masking view is created.

LUN addresses should only be supplied if the storage group is already contained within a view. In this case, a single LUN can be given, or one for each device range. If the LUN address is not specified, the array will assign the LUN address.

Restrictions

- Requires Storage Admin rights.
- Whether SG is FAST Managed or not, the command fails if adding FBA devices and the SG has CKD devices.
- Whether SG is FAST Managed or not, the command fails if adding CKD devices and the SG has FBA devices.
- Workload cannot be specified for CKD devices. The command fails if adding CKD devices to an SG if that SG has Workload set.

Remove devices from storage group

Description

Deleting a storage group requires that all devices are removed. However a storage group cannot be completely emptied if it is associated with a masking view,

Syntax

To remove devices or child storage groups from a storage group, use the following syntax:

```
symaccess -sid SymmID -name GroupName -type storage
  [-reserve_id ResvID [, ResvID [, ResvID...]]] [-force]
  [-unmap [-celerra]] [-ckd] [-rp]
remove devs SymDevStart:SymDevEnd | SymDevName |
  SymDevName, SymDevName, SymDevName...
remove -g DgName [-std] [-bcv] [-vdev] [-tgt]
remove -file DeviceFileName [src] [tgt]
remove sg SgName [, SgName1, SgName2, .., SgName]
```

Examples

For example, to remove the BCV devices in device group `Prog2` from storage group `SG_Prod` on array 458, enter:

```
symaccess -sid 458 -name SG_Prod -type storage remove -g Prog2 -bcv
```

Rename a storage group

Syntax

To rename a storage group, port group, or an initiator group, using the following syntax:

```
symaccess -sid SymmID rename -name GroupName -type <storage | port | initiator> -new_name
NewGroupName
```

Delete a storage group

Description

A storage group should be empty before it is deleted. It cannot be deleted if it is associated with a masking view or is in use by a FAST policy. To delete both a storage group and a masking view, delete the masking view first, then delete the storage group.

Syntax

To delete a storage group, use the following syntax:

```
symaccess -sid SymmIDdelete <view -name ViewName
  [-reserve_id ResvID[,ResvID[,ResvID...]]] > |
  < -name GroupName -type <storage
    [-reserve_id ResvID[,ResvID[,ResvID...]]] |
  port | initiator> [-force] > [-noprompt][-emulation ckd]
```

Options

-force

If the storage group still contains devices, forces the delete action.

Examples

To delete storage group SG_1 from array 458, enter:

```
symaccess -sid 458 -name sg_1 -type storage delete
```

List and show storage group information

Description

Storage group information is requested from the array or from the backup file. Information for a device range or a list of devices can be specified.

Syntax

To view storage group information, use the following syntax:

```
symaccess -sid <SymmID>
  [-offline] | -file <backup_filename>
  list -type storage
    [-devs <SymDevStart:SymDevEnd |
      SymDevName | <SymDevName,SymDevName...>]
    [-name GroupName] [-v | -detail]
  show GroupName -type storage
```

Manage 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. In addition, CHAP authentication can be enabled and disabled on port groups. This section details the following port management operations:

- Add ports
- Remove ports
- Delete port groups
- Copy a port group from array to array
- List and show port group information
- Lock down a Fibre Channel ID

NOTE: The `symaccess` control operations `add`, `remove`, `backup`, and `restore` are modified to support the extended FA director port range.

Add ports to port group

Description

Ports and iSCSI targets are added to an existing port group by specifying the name and type of the group, and the director port or iSCSI target information. Only a single front-end emulation of each type (FA, EF, etc.) can be assigned to each director, however each of these emulations can be assigned a variable number of physical ports (up to 32, numbered from 0 - 31). The `symaccess` control operations `copy -type port` or `copy view` copies the extended FA director port range (0 to 31) to the target array.

Syntax

To add ports to a port group, using the following syntax:

```
symaccess -sid SymmID -name GroupName -type port [-ckd][-celerra] [-rp]
  add -dirport Dir:Port [, Dir:Port [, Dir:Port...]]
  add -iscsi_tgt iSCSI [, iSCSI...]
```

Options

- `-ckd`
Specifies CKD devices.
- `-celerra`
Specifies Celerra devices.
- `-rp`
Specifies RecoverPoint devices.

Examples

To add port 4 of Fibre director 16D to port group PG_1 on array 245, enter:

```
symaccess -sid 245 -name PG_1 -type port add -dirport 16D:4
```

Remove ports from port group

Syntax

NOTE: A port group cannot be emptied if it is associated with a masking view. To remove a port or an iSCSI target, use the following syntax:

```
symaccess -sid SymmID -name GroupName
  -type port [-ckd][-force][-celerra] [-rp]
  remove -dirport Dir:Port [, Dir:Port [, Dir:Port...]]
  remove -iscsi_tgt iSCSI [, iSCSI...]
```

Options

- `-force`
Forces the port removal.

Example

To remove port 4 of Fibre director 16D from port group PG_1 on array 245, enter:

```
symaccess -sid 245 -name PG_1 -type port remove -dirport 16D:4
```

Delete port groups

Syntax

A port group cannot be deleted if it is associated with a masking view. To delete a port group, use the following syntax:

```
delete <view -name  
ViewName> [-force] | -name GroupName  
-type <storage | port | initiator > [-noprompt]
```

Options

-force

Forces the port group removal.

Examples

To delete port group PG_1 on array 245, enter:

```
symaccess -sid -name PG_1 -type port delete
```

Copy a port group from array to array

Syntax

To copy a port group from one array to another, use the following syntax:

```
symaccess -sid SymmID -target_sid SymmID copy -name GroupName -type port
```

List and show port group information

Syntax

To display port group information, use the following syntax:

```
symaccess -sid <  
SymmID [-offline] | -file <  
backup_filename>  
list -type port [-dirport Dir:Port]  
[-name GroupName] [-detail | -v]  
show GroupName -type port
```

Fibre Channel ID lockdown

Fibre Channel ID (FCID) lockdown is a security feature that limits host device access by adding Fibre Channel ID information of a switch within a fabric to device access records in the login history table. This feature handles WWN spoofing and the threat it poses to networked systems in a shared (same director port) storage port configuration.

This feature sets Fibre Channel ID (FCID) of the WWN of the HBA to be secured. The FCID is then added to the database record for the WWN of the specified HBA with the specified director and is locked. Once a Fibre Channel ID is locked, no user with a spoofed WWN can log in. If a user with a spoofed WWN is already logged in, that user loses all access through that HBA.

NOTE:

When an HBA logs in to a director port, the Fibre Channel ID accompanies it, indicating to the director port where to send its response. By specifying Fibre Channel ID information of the switch, the valid physical path through the SAN for a particular HBA is locked down. Only an HBA with a Fibre Channel ID that matches the FCID specified in the device masking record is able to log in to the storage port. It is recommended that at least two HBAs be available on the administrator host. If one HBA becomes locked out, the host will have access through the other HBA and can correct the record in the database.

Locking down a Fibre Channel ID

About this task

To find the Fibre Channel ID, lock it down, verify that it is locked down, and then force the change to take effect, use the following procedure:

Steps

1. Find the WWN. If the device is visible, run the `symaccess list hba` command to find the device path of the HBA to protect.
2. Find the Fibre Channel ID value.
3. Run the `symaccess set lockdown` command on with the FCID of the Fibre Channel ID found in step 2.

For example, to implement the Fibre Channel ID lockdown feature on Fibre Channel 021300 for director 16A, port 0, enter:

```
symaccess -sid 018 set lockdown on 021300
```

4. For the change to take effect, either reboot the host or pull the cable from the director and then replace the cable.

Manage initiator groups

This section describes how to manage initiator groups and includes the following tasks:

- Add initiators
- Remove initiators
- Delete initiators
- Initiator group flags
- Set HBA flags
- Replace a HBA
- Rename a HBA
- CHAP authentication

Add initiators to initiator group

Description

Initiators are added to an existing initiator group by specifying the initiator type (`-wwn` or `-iscsi`), the initiator group name, or by using an input file.

Syntax

To add initiators to an initiator group, use the following syntax:

```
symaccess -sid SymmID -name GroupName -type initiator  
-wwn wwn | -iscsi iscsi | -ig InitiatorGroupName |  
-f InitiatorFilename add
```

Add individual initiators to initiator group

Examples

To add initiator WWN 10000000c94ef69c to the initiator group IG_1 on array 245, enter:

```
symaccess -sid 245 -name IG_1 add -type initiator -wwn 10000000c94ef69c
```

Add initiators to initiator group using an input file

Examples

When using an input file, each initiator must be placed on a new line and start with either **WWN:** or **iSCSI:** or **IG:**, depending on the type of the initiator or initiator group name. The following is an example of the format for an initiator file:

```
WWN:10000000c94ef69c
iSCSI:iscsiname
IG:IgName
#WWN:10000000c94ef69d
```

NOTE: If the format of the initiator does not match the label at the start of the line, the file returns an error. A commented line, which the system ignores, is specified by placing the pound sign (#) at the beginning of a line.

Cascaded initiator groups

An initiator group can be added to another initiator group, only if it does not contain any initiator groups.

The following scenario describes cascaded initiator groups:

- HOST1 contains WWN1 & WWN2, which are added to IG_1.
- HOST2 contains WWN3 & WWN4, which are added to IG_2.
- IG_3 is created and contains IG_1 & IG_2.

In this example, gatekeeper devices for HOST1 can be assigned to IG_1, while different gatekeeper devices for HOST2 can be assigned to IG_2. The application devices needed by both hosts can be assigned to IG_3.

NOTE: If using the Volume Set Addressing flag, both the parent and child initiator group must have the flag.

Remove initiators from initiator group

Syntax

To remove an initiator from an initiator group, use the following syntax:

```
symaccess -sid SymmID -name GroupName -type initiator -wwn
wwn | -iscsi iscsi | -ig InitiatorGroupName
      | -f InitiatorFilename [-login]
      remove
```

Options

-login

Removes the initiator from the array's login history table.

Examples

To remove initiator WWN 10000000c94ef69c from the initiator group IG_1 on array 245, enter:

```
symaccess -sid 245 -name IG_1 remove -type initiator -wwn 10000000c94ef69c
```

Deleting initiator groups

Syntax

To delete an initiator group, use the following syntax:

```
delete <view -name  
ViewName> [-force] |  
-name  
GroupName -type <storage | port | initiator > [-noprompt]
```

Examples

To delete initiator group IG_1 on array 245, enter:

```
symaccess -sid -name IG_1 -type initiator delete
```

Initiator group flags

Table 20. Initiator group flags

Initiator group	Flag
Volume_Set_Addressing	[V]
Disable_Q_Reset_on_UA	[D]
Environ_Set	[E]
Avoid_Reset_Broadcast	[ARB]
OpenVMS	[OVMS]
SCSI_3	[SC3]
SPC2_Protocol_Version	[SPC2]
SCSI_Support1	[OS2007]

Set override flag for initiator group

Syntax

To set an override flag for an initiator group, use the following syntax:

```
symaccess -sid SymmID -name GroupName -type initiator  
set ig_flags <on <Flag> <-enable |-disable> | off [Flag]>
```

A flag cannot be set for the group if it conflicts with any initiator in the group. After a flag is set for a group, it cannot be changed on an individual initiator in the group.

Options

- on**
Turns on the specified initiator group port flag override and allows setting flag status of to *enabled* or *disabled*.
- off**
Turns off the specified initiator group port flag override.
- enable**
Sets the status of the initiator group port flag to *enabled*. The initiator group port flag override setting value must be *on* to set status.
- disable**

Sets the status of the initiator group port flag to *disabled*. The initiator group port flag override setting value must be on to set status.

Examples

To set the OS2009 [OS2009] flag for the initiator group `my_ig` on array 266, enter:

```
symaccess -sid 266 -type init -name my_ig set ig_flags on OS2009 -enable
```

To view the flag set in initiator group `my_ig` on array 266, enter:

```
symaccess -sid 266 show my_ig -type init -detail
```

Sample output

```
Symmetrix ID           : 000192600266
Initiator Group Name   : my_ig
Last updated at       : 10:52:15 AM on Wed Mar 31,2010
Port Flag Overrides   : Yes
  Enabled              : OS2009(OS2009)
                       Common_Serial_Number(C)
  Disabled            : Avoid_Reset_Broadcast(ARB)
Consistent Lun        : No
Originator Port wwn   : 1234567822446688
User-generated Name   : 1234567822446688/1234567822446688
FCID Lockdown         : No
Heterogeneous Host    : No
Port Flag Overrides   : Yes
  Enabled              : OS2009(OS2009)
                       Common_Serial_Number(C)
  Disabled            : Avoid_Reset_Broadcast(ARB)
CHAP Enabled          : N/A
Type                  : Fibre
```

Set HBA flags

Description

This feature allows specific host flags to be enabled and disabled on the director port. The HBA port flags are set on a per initiator basis, and the HBA must belong to an initiator group.

NOTE: Setting HBA port flags replaces setting the heterogeneous host configuration flags. To switch to setting HBA port flags, the heterogeneous host configuration must be disabled for a given HBA and all flags must be reset.

Syntax

To set (or reset) the HBA flags, use the following syntax:

```
symaccess -sid SymmID -wwn wwn | -iscsi iscsi
set hba_flags <on <flag,flag,flag...> <-enable |-disable> |
  off [flag,flag,flag...]>
```

Options

`hba_flags`

Sets the record in the database to hold information on the HBA port setting that may differ than the current setting on the FA.

`on | off`

Turns HBA flags on or off.

`flag`

Specifies the overridden HBA port flags as listed:

Table 21. Supported HBA flags

Supported HBA ports:	Supported initiator group ports:	Flag:
Avoid_Reset_Broadcast	Avoid_Reset_Broadcast	[ARB]
Disable_Q_Reset_on_UA	Disable_Q_Reset_on_UA	[D]
Environ_Set	Environ_Set	[E]
OpenVMS	OpenVMS	[OVMS]
SCSI_3	SCSI_3	[SC3]
SCSI_Support1	SCSI_Support1	[OS2007]
SPC2_Protocol_Version	SPC2_Protocol_Version	[SPC2]
Volume_Set_Addresssing	Volume_Set_Addresssing	[V]

-enable

Enables the specified HBA port flag(s) on a per initiator basis.

-disable

Disables the specified HBA port flag(s) on a per initiator basis.

Examples

To turn on HBA flags and enable the Common_Serial_Number and SCSI_3 flags, and disable the Disable_Q_Reset_on_UA flag on an HBA with the WWN 210000e08b0995b7 for array 031, director 16A port 0, enter:

```
symaccess -sid 031 set hba_flags on C,SC3 -enable -wwn 210000e08b0995b7-dir 16A -p 0
symaccess -sid 031 set hba_flags on D -disable -wwn 210000e08b0995b7-dir 16A -p 0
```

Sample output

The `symaccess show -detail` output displays the flags that are turned on and off for each HBA initiator that has the feature enabled.

```
symaccess -sid 237 -type initiator -detail show Prod1
Symmetrix ID          : 000190300237
Last updated at      : 08:46:54 AM on Tue Jul 29,2008
Initiator Group Name  : Prod1
Originator Port wwn   : 10000000c94ef69c
  User-generated Name : api196/10000000c94ef69c
  FCID Lockdown       : No
  Heterogeneous Host  : No
  Port Flag Overrides : No
  Type                 : Fibre
Originator Port wwn   : 5006016839a00c5c
  User-generated Name : 5006016839a00c5c/5006016839a00c5c
  FCID Lockdown       : No
  Heterogeneous Host  : No
  Port Flag Overrides : No
  Type                 : Fibre
iSCSI Name           : Symm_iScsi
  User-generated Name : iScsi_node_alias/iScsi_port_alias
  FCID Lockdown       : N/A
  Heterogeneous Host  : No
  Port Flag Overrides : No
  Type                 : iScsi
Group Name            : IniGrp
  User-generated Name : N/A
  FCID Lockdown       : N/A
  Heterogeneous Host  : N/A
  Port Flag Overrides : N/A
  Type                 : Initiator Group
```

Replacing a HBA

About this task

If a host adapter fails, or needs replacement for any reason, assign the devices associated with the old adapter to a new adapter use the `replace` action with the following syntax:

```
symaccess replace -wwn wwn -new_wwn NewWWN [-noprompt]
symaccess replace -iscsi iscsi -new_iscsi NewiSCSI [-noprompt]
```

To swap HBAs:

Steps

1. Run `symaccess list logins` to view the old WWN/iSCSI HBAs.
2. Swap the HBA boards according to the host instructions.
3. Run `symaccess list hba` or `discover` to view the new initiator (for example WWN).
4. Run `symaccess replace` to substitute a new WWN for all occurrences of the old WWN. For example, to replace old WWN 20000000c920b484 with new WWN 20000000c920b393:

```
symaccess -sid 814 replace -wwn 20000000c920b484 -new_wwn 20000000c920b393
```

5. Run `symaccess discover -rename` to establish the new AWWN and assign an AWWN to the new HBA in the login history table.

Rename a HBA

Syntax

To rename the alias for a specified initiator within a group, use the following syntax:

```
symaccess -sid SymmID rename -wwn wwn -alias alias
| -iscsi iscsi -alias alias
```

Using CHAP authentication

CHAP (Challenge Handshake Authentication Protocol) manages a credential name and a CHAP secret, which are similar to a username and a password, though more secure than the standard Password Authentication Procedure (PAP).

Enable CHAP on an iSCSI initiator

Syntax

To enable CHAP on an iSCSI initiator, use the following syntax:

```
symaccess -sid SymmID -iscsi iscsi enable chap
```

Enable CHAP on a director and port

Syntax

To enable CHAP on a director and port, use the following syntax:

```
symaccess -sid SymmID [-dirport Dir:Port]enable chap
```

Set the CHAP credential and secret

Syntax

To set the CHAP credential and secret, use the following syntax:

```
symaccess -sid SymmID -dirport Dir:Port | -iscsi iscsi  
set chap -cred Credential -secret Secret
```

Disable CHAP on a specific director and port

Syntax

To disable CHAP on a specific director and port, use the following syntax:

```
symaccess -sid SymmID [-dirport Dir:Port] disable chap
```

Delete CHAP from a specific director and port

Syntax

To delete CHAP from a specific director and port, use the following syntax:

```
symaccess -sid SymmID [-dirport Dir:Port] delete chap
```

Display CHAP information

Examples

To display CHAP information for array 001.

```
symaccess list chap -sid 001
```

Sample output

```
Symmetrix ID           : 000197100001  
Director Identification : SE-9G  
Director Port          : 310  
iSCSI Target Name     : iqn.1992-04.com.emc:sn.11121318  
Protocol               : CHAP  
Identifier             : Type      State      Credential  
-----  
SE-9G:310              : N/A      ENABLED   credential  
symaccess list chap -sid 001 -v  
Symmetrix ID           : 000197100001  
Director Identification : SE-9G  
Director Port          : 310  
iSCSI Target Name     : iqn.1992-04.com.emc:sn.11121318  
Protocol               : CHAP  
Identifier             : SE-9G:310  
Identifier Type        : N/A  
State                  : ENABLED  
Last updated at       : 09:58:27 AM on Fri Apr 17, 2015  
CHAP Credential       : credential
```


Verify the auto-provisioning database

Description

Use the `symaccess verify` command to verify that the auto-provisioning database is consistent. Any inconsistencies display in the command output. This command can also be used with a backup file. Add the `-log` option for reporting the inconsistencies in a log file.

Syntax

To verify the auto-provisioning database, use the following syntax:

```
symaccess -sid SymmID | -file BackupFileName verify [-log]
```

Options

-log

Reports database inconsistencies in a log file

Examples

When database is consistent:

```
symaccess -sid 266 verify
```

```
Starting a verify operation.....  
The auto provisioning database is consistent
```

Verification of a database backup file when database has inconsistencies:

```
symaccess -file /tmp/bkup1.file verify
```

```
Starting a verify operation.....  
Found SG 'stor_grp1' to contain the view flag but didn't find a matching view  
Found IG 'init_grp1' contains invalid initiator records  
Found masking view 'mask_view1' with parent IG 'init_grp1' but no masking records for the  
child IG 'child_grp1' are present  
There are inconsistencies in the auto provisioning database
```

Display auto-provisioning group information

Syntax

To display auto-provisioning group information, use the following syntax:

```
symaccess -sid SymmID [-offline] | -file BackupFilename list [-name GroupName] [-v]  
list -type <storage [-devs <SymDevName:SymDevName>] | port  
[-dirport Dir:Port] | initiator [-wwn wwn |  
-iscsi iscsi> [-name GroupName] [-detail | -v]  
list devinfo [-ig InitiatorGroupName]  
list view [-name ViewName] [-v] [-detail]  
show GroupName -type <initiator [-detail] | port | storage>  
show view ViewName [-ig ChildInitiatorGroupName]  
symaccess -sid SymmID | -file BackupFilename list chap [-dirport Dir:Port] [-v]  
symaccess -sid SymmID list assignment [-v] -devs <SymDevStart:SymDevEnd |  
SymDevName | <SymDevName,SymDevName...>  
list no_assignments [-dirport Dir:Port
```

Options

-offline

For use with the `symaccess list` and `show` commands. When the `-offline` option is used, the command reports the data from the `symapi` configuration database file and not from the array. Alternatively, set the `SYMCLI_OFFLINE` environment variable to 1 to enable the offline mode for reporting.

`-detail`

Displays all auto-provisioning details. Without using the `-detail` option, any column without data does not display. If the `-detail` option is provided, the column without data displays a dash (-).

Modified reporting commands for extended FA director port range

Only a single front-end emulation of each type (FA, EF, etc.) can be assigned to each director, however each of these emulations can be assigned a variable number of physical ports (up to 32, numbered from 0 - 31). The following `symaccess` reporting commands support the extended FA director port range:

- For `symaccess show -type port,show view`, and `symaccess show backupfile view` output, the Director Identification section supports reporting for the extended port range. In addition new fields are added: the `WWN port name`, which displays the Fibre channel director port WWN, and the `iSCSI name`, which displays the name of the iSCSI director port.
- For `symaccess list hba`, `symaccess list devinfo` output, and `symaccess list assignments` the `Dir:Port` field supports reporting for the extended port range.
- The `symaccess list chap` output supports reporting for the extended port range. In addition, a new field is added, the `iSCSI Target Name`, which displays the iSCSI name for the iSCSI director port.
- For `symaccess list no_assignments` and `symaccess list logins` output, the `Director Port` field supports reporting for the extended port range.

NOTE:

The `symaccess list no_assignments` output reports that "All devices available on this director/port are assigned" for ACLX enabled ports.

Display masking views

Description

The `symaccess show view` command lists the masking view and the associated initiator group, the port group, the storage group, and any child groups. This command also displays the time the group was modified and the time the associated masking view was modified.

Options

`-detail`

Displays all masking view details, and columns without data displays a dash (-). Without using the `-detail` option, any column without data does not display.

Examples

To list all masking views on array 237, enter:

```
symaccess -sid 237 list view
```

To list the details for masking view `view1_86`, enter:

```
symaccess -sid 0225 show view view1_86 -detail
```

Use the `-detail` option with the `show` command to display all the details of a view, including the child initiator groups.

Sample output

For all masking views on array 237:

```
Symmetrix ID           : 000190300237
Masking View Name      Initiator Group Port Group Storage Group
-----
View1                  IG_1          PG_1      SG_1
View2                  WinHost       PG
...
```

Detailed output for masking view view1_86 :

```
Symmetrix ID           : 000196700255

Masking View Name      : view1_86
Last update time       : 04:36:39 AM on Mon Aug 31,2015
View last update time  : 04:36:39 AM on Mon Aug 31,2015

Initiator Group Name   : ig1_86

  Host Initiators
  {
    IG   : ig2_86
  }
Port Group Name        : pg1_86

  Director Identification
  {
    Director
    Ident Port   WWN Port Name / iSCSI Target Name
    -----
    FA-1D  027 500009735003fc1b
  }
Storage Group Name     : sg2_86

  Number of Storage Groups : 0
  Storage Group Names     : None

Masking View Name      : view1_86
Last update time       : 04:36:39 AM on Mon Aug 31,2015
View last update time  : 04:36:39 AM on Mon Aug 31,2015

Initiator Group Name   : ig2_86 *

  Host Initiators
  {
    WWN   : addddeb9db4bf88f
          [alias: addddeb9db4bf88f/addddeb9db4bf88f]
  }
Port Group Name        : pg1_86

  Director Identification
  {
    Director
    Ident Port   WWN Port Name / iSCSI Target Name
    -----
    FA-1D  027 500009735003fc1b
  }
Storage Group Name     : sg2_86

  Number of Storage Groups : 0
  Storage Group Names     : None

Sym   Dir:Port   Physical Device Name   Host   Attr   Cap(MB)
----  -
00060  01D:027   Not Visible             1      -----  1796
-----

Total Capacity                1796
```

* Denotes a cascaded Initiator Group within the specified Masking View

View auto-provisioning group details

Syntax

To display storage group, port group, or initiator group details, using the `symaccess list` and `symaccess show` commands, use the following syntax:

```
symaccess -sid SymmID [-offline] | -f BackupFilename
  list [-name GroupName] [-v]
  list -type <storage [-devs SymDevName <:SymDevName>] | port
    [-dirport Dir:Port] | initiator [-wwn wwn |
    -iscsi iscsi]> [-v] [-name GroupName]

  show GroupName -type <storage | port | initiator [-detail]>
  show view ViewName
```

Options

- type**
Use this option to list groups of a specific type: (initiator, port, or storage).
- v**
Lists details of a group and any children in the group.
- details**
Shows detail of a masking view or any of the groups in the masking view.

Examples

To list details of a storage group or any child storage group, enter:

```
symaccess list -type storage -v
```

To show details of a masking view or any of the groups in the masking view, enter:

```
symaccess show view -detail
```

Sample output

For storage group:

```
Symmetrix ID           : 000195700601
Storage Group Name     : backup_storage
Device Count           : 1
Storage Group Count    : 0
Masking View Count     : 0
Last update time      : 12:32:26 AM on Fri Apr 06,2012
Group last update time : 12:32:26 AM on Fri Apr 06,2012
Masking View Names     : None
Storage Group Name     : Mkt_storage (IsParent)
Device Count           : 0
Storage Group Count    : 1
Masking View Count     : 0
Last update time      : 08:10:22 PM on Tue Feb 14,2012
Group last update time : 10:32:26 AM on Fri Apr 06,2012
Masking View Names     : Mkt_view *
. . .
* Denotes Masking Views through a cascaded group
```

For masking view:

```
Symmetrix ID          : 000195700601
Masking View Name     : HR_view
Last update time      : 04:30:24 AM on Tue Apr 03,2012
View last update time : 04:35:36 AM on Tue Apr 03,2012
Initiator Group Name  : HR_hosts
  Host Initiators
  {
    IG   : HR_host_GrpA
  }
Port Group Name       : HR_ports
  Director Identification
  {
    FA-15E:1
  }
Storage Group Name    : HR_storage
  Number of Storage Groups : 0
  Storage Group Names   : None
Sym
Dev      Dir:P  Physical Device Name      Host  Attr  Cap(MB)
-----
01C27   15E:1  Not Visible                          1     ----  300000
-----
Total Capacity                               300000
```

View auto-provisioning device assignments

Syntax

To display the assignments for one or more devices, use the following syntax:

```
symaccess -sid
SymmID list assignments
-devs
SymDevStart:
SymDevEnd |
SymDevName |
   SymDevName,
SymDevName... [-v]
```

Examples

To list the assignments for device range 20:22 and device 24 on array 120, enter:

```
symaccess -sid 120 list assignments -devs 20:22,24
```

Sample output

```
Symmetrix ID          : 000192600120
Device  Identifier      Type  Dir:P
-----
00020   10000000c9594dce  FIBRE FA-7E:1
        210000e08b04daac FIBRE FA-7E:1
        210000e08b1ed7f1 FIBRE FA-7E:1
00021   10000000c9594dce  FIBRE FA-7E:1
        210000e08b04daac FIBRE FA-7E:1
        210000e08b1ed7f1 FIBRE FA-7E:1
00022   10000000c9594dce  FIBRE FA-7E:1
        210000e08b04daac FIBRE FA-7E:1
        210000e08b1ed7f1 FIBRE FA-7E:1
00024   10000000c9594dce  FIBRE FA-7E:1
        210000e08b04daac FIBRE FA-7E:1
        210000e08b1ed7f1 FIBRE FA-7E:1
```

Display device list with no auto-provisioning assignments

Syntax

To display a list of devices with no assignments, use the following syntax:

```
symaccess -sid SymmID list no_assignments  
[-dirport DirNum:PortNum]
```

Examples

To list the devices without assignments for array 120, enter:

```
symaccess -sid 120 list no_assignments
```

Sample output

```
Symmetrix ID           : 000192600120  
Director Identification : FA-7F  
Director Port          : 0  
ACLX Enabled           : No  
No devices were found for this director/port  
Director Identification : FA-7F  
Director Port          : 1  
ACLX Enabled           : Yes  
Devices not yet assigned :  
00030  
00031  
00032  
00033  
00034  
...
```

List initiator group devices

Syntax

To list all the devices masked to an initiator group, use the following syntax:

```
symaccess -sid SymmID list devinfo [-ig InitiatorGroupName]
```

When using the `-detail` option, the initiator group displays include the child initiator group device information.

View the HBA alias name

Example

To display the alias names for masking view `host1082_view`, enter:

```
symaccess show host1082_view -sid 001 view
```

Sample output

Display supports up to 32 characters each for the alias node name and the alias port name.

```
Symmetrix ID           : 000197100001  
Masking View Name      : host1082_view  
Last updated at        : 12:51:37 PM on Tue May 13, 2014  
View last update time  : 01:37:41 PM on Thu Apr 02, 2015
```

```

Initiator Group Name      : hba1_2_3
Host Initiators
{
  WWN      : 10000000C99DE136
            [alias: api019010000000C99DE136/api019010000000C99DE136]
  ISCSI    : iqn.2002-06.com.host1082
            [alias: api1082/api1082]
}
Port Group Name          : host1082_ports
Director Identification
{
  Director
  Ident  Port  WWN Port Name / iSCSI Target Name
  -----
  FA-3E  001  5000097300092110
  SE-9H  000  iqn.1992-04.com.emc:5000097300092190
  SE-10G 000  iqn.1992-04.com.emc:5000097300092194
}
Storage Group Name       : application_sg
Number of Storage Groups : 0
Storage Group Names      : None
Sym
Dev  Dir:Port  Physical Device Name  Host  Lun  Attr  Cap (MB)
-----
00083  3E:001  Not Visible           1     1     6
        9H:000  Not Visible           1     1     6
        10G:000  Not Visible           1     1     6
-----
Total Capacity           6

```



Querying and Reporting

This section describes the SYMCLI query commands and various tools for monitoring and managing storage array performance.

Chapters include:

Topics:

- [Configuration Query Operations](#)
- [Events and Logs](#)
- [XML Structured Output](#)

Configuration Query Operations

This chapter describes how to use the SYMCLI to collect and display configuration data for array disks, and virtual environments.

Topics:

- [Configuration data overview](#)
- [SCSI-level data](#)
- [Array-level data](#)
- [Device-level data](#)
- [Disk-level data](#)

Configuration data overview

The Solutions Enabler SYMCLI provides various commands that are used to query different levels of the storage environment. These levels include:

- **SCSI-level** — Returns data at the SCSI level, whereby the SYMCLI issues SCSI INQUIRY and SCSI READ CAPACITY to return low-level physical device data (such as vendor, configuration, and basic configuration) and host HBA information (such as vendor, model, firmware, and basic configuration).
- **Array-level** — Returns detailed data about the configuration of one or all arrays. The data returned includes host relationship information (local or remote), cache size, and number of devices. Other details include vendor specific configuration information, ports, flags, adapters, Data at Rest Encryption data, pool information, and environment data.
- **Device-level** — Returns device level data on locally or remotely attached arrays. Device-level data includes capacity, cache, emulation, configuration, group, and usage information that is critical to all other storage management operations.
- **Disk-level** — Returns detailed information about the raw disks such as detailed identification, vendor, and capacity.
- **Virtual environment** — Returns configuration information about a VMware or Microsoft Hyper-V environment, such as the type and capacity of their storage pools.

SCSI-level data

The `syminq` command is used to obtain SCSI disk device information using a SCSI INQUIRY command, and optionally SCSI READ CAPACITY command, on host HBAs and one or all locally attached physical devices. This command returns SCSI-level data for EMC arrays, StorageWorks, or HDS devices.

The `syminq` command can return a list all of the HBAs on the local host. Using the options available, the scope of this request can be limited to only the Fibre HBAs, SCSI HBAs, iSCSI HBAs, or HBAs derived from a SNIA API query. In the UNIX environment, if the SNIA libraries are unavailable, HBA information is obtained from the lost log files.

Returned results can be limited to devices with target mapping information (only devices mapped through fibre HBAs using the `-mapinfo` option), or to device identifiers (by a user or application by using the `-identifier` option). Identifiers are limited to `device_name`, `nice_name`, `hp_id`, and `vms_id`.

The *EMC Solutions Enabler SYMCLI Command Reference Guide* describes the various options available with the command.

NOTE:

Data returned from issuing any `syminq` command is not stored in the configuration database.

SCSI device list

Examples

To list all devices listed by physical device name, enter:

```
syminq
```

Sample output

Device		Product			Device		
Name	Type	Vendor	ID	Rev	Ser Num	Cap (KB)	fsy
/dev/sda		MAXTOR	ATLAS10K4_3*	DFL0	B2CS5FSM	35916548	
/dev/sdb	GK	EMC	SYMMETRIX	5977	150009E020	5760	
/dev/sdi	GK	EMC	SYMMETRIX	5977	15000A7020	5760	
/dev/sdj	GK	EMC	SYMMETRIX	5977	15000A8020	5760	
/dev/sdk	GK	EMC	SYMMETRIX	5977	15000A9020	5760	
/dev/sdl	GK	EMC	SYMMETRIX	5977	15000AA020	5760	
/dev/sdm	GK	EMC	SYMMETRIX	5977	15000AB020	5760	
/dev/sdn	GK	EMC	SYMMETRIX	5977	0700022000	5760	
/dev/sdo	GK	EMC	SYMMETRIX	5977	0700023000	5760	
/dev/sdp	GK	EMC	SYMMETRIX	5977	0700024000	5760	
/dev/sdq	GK	EMC	SYMMETRIX	5977	0700025000	5760	
. . .							

Devices listed by array ID

Examples

To list devices by array ID, enter:

```
syminq -symmids
```

Sample output

Device		Symmetrix		Device		
Name	Type	ID	Rev	Ser Num	Cap (KB)	
/dev/sdb	GK	000190300215	5977	150009E020	5760	
/dev/sdc	GK	000190300215	5977	150009F020	5760	
/dev/sdd	GK	000190300215	5977	15000A2020	5760	
/dev/sde	GK	000190300215	5977	15000A3020	5760	
/dev/sdf	GK	000190300215	5977	15000A4020	5760	
/dev/sdg	GK	000190300215	5977	15000A5020	5760	
/dev/sdh	GK	000190300215	5977	15000A6020	5760	
. . .						

List of devices without capacity

Examples

To list devices without issuing a SCSI READ CAPACITY, enter:

```
syminq -symmids -nocap
```

Sample output

Device		Symmetrix		Device	
Name	Type	ID	Rev	Ser	Num
/dev/sdv	BCV	000190300516	5977	1600028000	
/dev/sdw	BCV	000190300516	5977	1600029000	
/dev/sdx	BCV	000190300516	5977	160002A000	
/dev/sdy	BCV	000190300516	5977	160002B000	
/dev/sdz	BCV	000190300516	5977	160002C000	

List of devices with WWN

Options

-colon

Use this option to return a list of devices with WWN for each device using colons as spacers.

Examples

To list the WWN for each device, enter:

```
syminq -wwn
```

Sample output

Device		Device		
Name	Num	Array ID	WWN	
/dev/sda	N/A	N/A	N/A	
/dev/sdb	0005E	000194900306	60000970000194900306533030303545	
/dev/sdy	000E6	000194900306	60000970000194900306533030304536	

Output with **-colon** option:

Device		Device		
Name	Num	Array ID	WWN	
/dev/sda	N/A	N/A	N/A	
/dev/sdb	0005E	000194900306	60:00:09:70:00:01:94:90:03:06:53:30:30:30:35:45	
/dev/sdy	000E6	000194900306	60:00:09:70:00:01:94:90:03:06:53:30:30:30:45:36	

Device list in pdevfile format

Examples

To list device names in a format suitable for use as pdevfile:

```
syminq -pdevfile
```

Sample output

#	Symm_id	pdev	dev	dir	dir_port
000187900771	/dev/rdisk/c2t0d0s2	0017	15D	0	
000187900771	/dev/rdisk/c2t0d1s2	0018	15D	0	

List device mapping information

Syntax

To list mapping information, use the following syntax:

```
syminq -mapinfo [PdevName]  
               [-sym[-powerpath] | -hds | -storworks]  
               [-cache | -nocache] [-colons] [-winvol]
```

Examples

```
syminq -mapinfo
```

Sample output

Device	Target Mapping		
Name	HBA Port	WWN	Target Port WWN
/dev/rdsd/c0t0d0s2	N/A		N/A
/dev/rdsd/c1t1d0s2	N/A		N/A
/dev/rdsd/c2t5000097208013558d0s2	210000e08b80a234		5000097208013558
/dev/rdsd/c2t5000097208013558d1s2	210000e08b80a234		5000097208013558
/dev/rdsd/c2t5000097208013558d2s2	210000e08b80a234		5000097208013558
/dev/rdsd/c2t5000097208013558d3s2	210000e08b80a234		5000097208013558
/dev/rdsd/c2t5000097208013558d4s2	210000e08b80a234		5000097208013558

List device identifiers

Description

To list the array device identifiers assigned to devices by the user or other applications, use the `syminq -identifier` option and specify one of the four options: `device_name`, `nice_name`, `hp_id`, or `vms_id`.

Examples

To list the VMS device identifiers, enter:

```
syminq -identifier vms_id
```

Sample output

Device			Device	
Name	Num	Vendor	Array ID	VMS ID
/dev/sdb	0005E	EMC	000194900306	94
/dev/sdc	0005F	EMC	000194900306	95
/dev/sdy	000E6	EMC	000194900306	230

Array-level data

The `symcfg` command provides arguments and options to report specific array-level components. The storage array data includes the following:

- Storage arrays — A list of all arrays attached to the host and configuration details about them.

- Registered applications — Applications registered with SYMAPI that have accessed all or specified arrays to which your host is connected.
- Host connections — Detailed information about the hosts that have accessed a array, the host node name and the array ID, director/port information, and array capacity.
- PowerPath host registration — Detailed information about the PowerPath hosts that have registered on an array, including the host name, array ID, OS version, PowerPath version and cluster information.
- Director information — Configuration and status information about all directors of a specified array including address, port, and status information. Details can be limited to a specific type of director.
- SELs (Symmetrix External Locks) and Semaphores — Status information on SELs (by number or type) and SYMAPI semaphores.
- Cache management — Information on the LRU cache, including the cache slots that each LRU occupies, and the percentage of the total cache utilized.
- Network services — A list of network services available to SYMAPI client applications.
- Environment data — Detailed information on the memory boards, and the array's environmental data, including fans and power supplies, can be obtained.
- CU Images — Information for mainframe users.
- Microcode patches — A list of installed operating environment for array patches.
- Bay location descriptions — A list of the bay names and location details of the array.
- Data at Rest Encryption — Information about hardware-based, on-array, back-end encryption for arrays.

List arrays

Description

To query storage environment configuration data, the arrays in the storage environment must first be identified. Each array has a serial number (SID) that is used to uniquely identify it. The `symcfg` command lists of all the accessible arrays by SID including the model number and the number of accessible devices. Use the SID with other command options to obtain configuration information for the array directors and devices.

Examples

To list all array IDs connected to a host, enter:

```
symcfg list
```

To obtain a detailed list for a specific array, using the verbose option (`-v`), enter:

```
symcfg list -v -sid 2300
```

Sample output

Output listing all arrays connected to a host.

S Y M M E T R I X						
SymmID	Attachment	Model	Mcode Version	Cache Size (MB)	Num Phys Devices	Num Symm Devices
000190100097	Local	VMAX40K	5977	98304	20	32088
000000006206	Remote	VMAX40K	5977	32768	0	4139
000187900035	Remote	VMAX40K	5977	8192	0	1314
000187900041	Remote	VMAX40K	5977	8192	0	960
000190300174	Remote	VMAX40K	5977	16384	0	335
000190300175	Remote	VMAX40K	5977	16384	0	914
000190300184	Remote	VMAX40K	5977	16384	0	417

Output listing details (`-v`) for array 2300.

```
Product Model           : VMAX40K
Symmetrix ID            : 000195602300
Microcode Version (Number) : 5977 (16F40000)
Microcode Registered Build : 0
Microcode Date          : 10.16.2012
Microcode Patch Date    : 10.16.2012
Microcode Patch Level   : 159
```

```

Symmwin Version           : 30
Engenuity Build Version  : 5977.159.30
Service Processor Time Offset : - 00:00:02
Cache Size (Mirrored)     : 18432 (MB)
# of Available Cache Slots : 253968
Max # of System Write Pending Slots : 152777
Max # of DA Write Pending Slots : 0
Max # of Device Write Pending Slots : 7638
Max # of Replication Cache Slots : 43174
Replication Usage (Percent) : 10

```

. . .

Reporting differences

The following items describe reporting differences that are dependent on HYPERMAX, and Solutions Enabler versions:

- For HYPERMAX OS 5977 lower than Q12016SR, the `Replication Usage (Percent)` is reported as N/A.
- For compatibility modes earlier than 8.2 the `Replication Usage (Percent)` field displays as `Replication Cache Usage (Percent)`.

List arrays by system resource demand

Description

The `symcfg list -demand` command reports system resource demand. This feature is supported on arrays running HYPERMAX OS 5977 Q22017SR or higher.

Examples

To list system resource demand, enter:

```
symcfg list -demand -sid 188
```

To list resource demand details for array 188, enter:

```
symcfg list -demand -v -sid 188
```

To list metadata details for array 188, enter:

```
symcfg list -demand -md -sid 188
```

To list array capacity details for array 188, enter:

```
symcfg list -demand -detail -sid 188
```

Sample output

```
symcfg list -demand -sid 188
```

```
A R R A Y      C A P A C I T Y      R E P O R T
```

```

-----
                Subscribed Capacity  Snapshot Capacity  Usable Capacity
                -----
                Total  Allc          Total Mdfy      Total  Used
SymmID          (TB)  (%)          (TB)  (%)      (TB)  (%)
-----
000197800188   208.00   77          416.00  14      116.30   66

```

Output with `-verbose (-v)` option for specified array:

```
symcfg list -demand -v -sid 188
```

```

Symmetrix ID           : 000197800188
Array Usage
  User Subscribed (TB)  : 208.00
  Allocation (%)        : 77
  Non Shared (TB)      : 144.00
  Shared (TB)          : 16.00
  System Subscribed (TB) : 2.00
  Snapshot capacity (TB) : 416.00
  Modified (%)         : 14
  Non Shared (TB)      : 41.60
  Shared (TB)          : 14.60
  Usable Capacity (TB) : 116.30
  Used (%)             : 66
  Used Capacity (TB)   : 76.20
  User Used (TB)       : 74.20
  System Used (TB)     : 2.00
  Temp Used (TB)       : 22.30
  Encapsulated (TB)    : 0.00
Meta Data Usage
  System (%)           : 77
  Replication (%)      : 95
  Frontend (%)         : 70
  Backend (%)          : 75

```

Metadata option (-md) for specified array:

```
symcfg list -demand -md -sid 188
```

```
M E T A D A T A   U S A G E
```

```

-----
SymmID           System  Replication  Frontend  Backend
                Used (%)  Used (%)    Used (%)  Used (%)
-----
000197800188    77          95          70        75

```

```
symcfg list -demand -detail -sid 188
```

```

                A R R A Y   C A P A C I T Y   R E P O R T
                Subscribed Capacity          Snapshot Capacity          Usable Capacity
-----
SymmID           Total Allc NonShared Shared  Total Mdfy NonShared Shared  Total Used  User System Temp
                (TB) (%) (TB) (TB) (TB) (%) (TB) (TB) (TB) (%) (TB) (TB) (TB)
-----
000197800188    208.00  77  144.00  16.00  416.00  14  41.60  14.60  116.30  66  76.20  2.00  22.30

```

Data at Rest Encryption

Data at Rest Encryption (D@RE) provides hardware-based, on-array, back-end encryption for VMAX family arrays. Back-end encryption protects information from unauthorized access when disk drives are removed from the system. Data Encryption provides encryption on the back end using Fibre Channel I/O modules that incorporate AES-XTS 256-bit data-at-rest encryption. These modules encrypt and decrypt data as it is being written to or read from disk. All configured drives are encrypted, including data drives, spares, and drives with no provisioned volumes. In addition, all array disk data is encrypted, including array File System and PowerVault contents.

Data at Rest Encryption supports either an internal embedded key manager, or RSA Data Protection Manager for external, enterprise-grade key management. For external key management, Data at Rest Encryption is qualified for interoperability with the RSA Key Manager Appliance version 2.7 SP1, and the RSA Data Protection Manager (DPM) version 3.1 (appliance).

 **NOTE: For more information refer to the *EMC VMAX3 Family Product Guide*.**

By securing data on enterprise storage, Data Encryption ensures that the potential exposure of sensitive data on discarded, re-used, or stolen media is reduced or eliminated. As long as the key used to encrypt the data is secured, encrypted data cannot be read. In addition to protecting against threats related to physical removal of media, this also means that media can readily be repurposed by destroying the encryption key used for securing the data previously stored on that media. In this way, disk rotation, migration and upgrade are secured, without changes to operational procedures.

Data Encryption is compatible with all array system features, allows for encryption of any supported local drive types or volume emulations, and provides encryption without performance degradation or disruption to existing applications or infrastructure.

NOTE: All key management is transparent to the storage administrator; no direct control of keys is allowed through Solutions Enabler. Encryption must be turned on by an EMC Customer Engineer during the installation of the array. There is no way to turn encryption on or off using Solutions Enabler.

List data encryption status

Options

-v

Lists array details including data encryption status.

Examples

To list data encryption status (using -v option) for array 343, enter:

```
symcfg list -sid 343 -v
```

Sample output

If encryption has not been turned on at install time, the `symcfg -v` command displays Symmetrix Data Encryption as Disabled.

If encryption is not supported on the array, the `symcfg -v` command displays Symmetrix Data Encryption as N/A.

```
Symmetrix ID: 000194900343
Time Zone   : Eastern Standard Time
Product Model      : VMAX
Symmetrix ID      : 000194900343
. . .
3 Dynamic Mirrors      : Enabled
Cache Partitioning    : Disabled
IPSec Status         : Pass Thru
Allow spare in mirror 4 position : Disabled
Disks Service        : Customer Replaceable
Symmetrix Data Encryption : Enabled
. . .
```

List Data encryption key (DEK) audit log

Description

Every key management event generates an entry in the array audit log, including initial DEK configuration information.

Example

```
symaudit list -text -sid 012 -function_class Security
```

Sample output

```
A U D I T   L O G   D A T A
Symmetrix ID      : *****012
Record           :
Number          Date          Time          Function Class      Action Code      Activity ID
-----
Text
-----
2 11/30/10 14:50:55 Security  Init      SW_Men_6829
DARE Local Install.
3 11/30/10 14:50:55 Security  Create   SW_Men_6829
New KEK key generated. MUID:
61A3770EDBEE3154463E15A5FA91006C6DA4F9C14C96653635962DE21F636125
4 11/30/10 14:50:55 Security  Create   SW_Men_6829
New DEK key generated. Drive WWN: 20000024b65d871b Location: DA07 A0:00 MUID:
828547DA8EA4D623F674C1FE659329A0D8487308E7BC77979BABC8554ED9B7BF
5 11/30/10 14:50:55 Security  Create   SW_Men_6829
```

```

New DEK key generated. Drive WWN: 20000024b65d3606 Location: DA07 A0:02 MUID:
C0D587D1EEBC66196EFB63449E31281B50AB3F6B048440661334056FF18EFD4C
 6 11/30/10 14:50:55 Security Create SW_Men_6829
New DEK key generated. Drive WWN: 20000024b65ddb94 Location: DA07 A0:04 MUID:
2D518DC19DB2D6856DEC5064653EC59CE8FAEA36D11A2078ECA088589F62A647
 7 11/30/10 14:50:55 Security Create SW_Men_6829
New DEK key generated. Drive WWN: 20000024b65d9bef Location: DA07 A0:06 MUID:
3738B80DF56F65F3E1526E7F16A2721ACED2774A08A35A8B7ABF1442C665F2F1
 8 11/30/10 14:50:55 Security Create SW_Men_6829
New DEK key generated. Drive WWN: 20000024b65ddb3c Location: DA07 A0:08 MUID:
40932F670A3CF4B15909EA309D3A10E1CDD66ED4D5564F75BD75F96A9CA58040
 9 11/30/10 14:50:55 Security Create SW_Men_6829
New DEK key generated. Drive WWN: 20000024b6864000 Location: DA07 A0:0A MUID:
99234F816501718AA87B29C525976DF7EE9351DC6CAD5756DB7347345F27E3BB
10 11/30/10 14:50:55 Security Create SW_Men_6829
New DEK key generated. Drive WWN: 20000024b64e634e Location: DA07 A0:10 MUID:
0FC78B4ECA29

```

Post-drive replacement DEK audit log

Description

When a drive replacement results in the destruction of the Data Encryption Key for the original drive and the creation of a DEK for the new drive, these events are logged in the array Audit Log.

Example

The following output lists this information for array 012:

```
symcaudit list -text -function_class Security -sid 012
```

Sample output

```

A U D I T   L O G   D A T A
Symmetrix ID : *****012
Record      Function  Action  Activity
Number     Date      Time    Class   Code    ID
-----
Text
-----
135 11/30/10 10:15:58 Security Modify  SW_Men_E757
      DEK key deactivated. Drive WWN: 20000024b654edce Location: DA08 C1:3E MUID:
E0252F83955C803D03596CF3B3331CED715FA985A6F99E2F6EC0DF325A47458E
136 11/30/10 10:15:59 Security Delete  SW_Men_E757
      DEK key destroyed. Drive WWN: 20000024b654edce Location: DA08 C1:3E MUID:
E0252F83955C803D03596CF3B3331CED715FA985A6F99E2F6EC0DF325A47458E
142 11/30/10 10:17:49 Security Create   SW_Men_E757
      New DEK key generated. Drive WWN: 20000024b65dd7ea Location: DA08 C1:3E MUID:
5EC4CF4E2B4D92C98A9FE9A8A19E527B090D95E75C305DAD21E98ED630C5A467

```

View application registrations with array access

Examples

For example, to list all the applications that have host to array connection for array 282, enter:

 **NOTE: Omit the `-sid` option list connections for all arrays.**

```
symcfg list -applications -sid 282
```

Sample output

NOTE: If Embedded Management is configured on the VMAX array, the node name and IP address listed in the output is the nodename and IP address identified by the NAT gateway, and not the internal identity of the eManagement Guest. For more information on eManagement refer to the *EMC VMAX3 Family Product Guide*.

```
Symmetrix ID      : 000192600282
      Host
```

Host		Application			
Node Name	IP Address	ID	Vendor ID	Version	Attr
HK192600282	0.0.0.1	SYMACCESS	EMC Corp	7.3.0.194	-
		EVTdaemon	EMC Corp	7.3.0.194	-
		SMBASE	EMC Corp	7.3.12.9	-
		SYMACCESS	EMC Corp	7.3.0.177	-
		EVTdaemon	EMC Corp	7.3.0.177	-
		SMBASE	EMC Corp	7.3.11.8	-
		EVTdaemon	EMC Corp	7.3.0.175	-
		SYMLMF	EMC Corp	7.3.1202.0	-
		SYMACCESS	EMC Corp	7.3.0.175	-
api101	10.247.80.101	SYMCFG	EMC Corp	7.3.0.234	(R)
		SYMRDF	EMC Corp	7.3.0.230	(M)
		SYMCG	EMC Corp	7.3.0.230	(M)
		GNSdaemon	EMC Corp	7.3.0.230	(R)
		SYMCFG	EMC Corp	7.3.0.230	(R)
		SYMCG	EMC Corp	7.3.0.227	(M)
		SYMRDF	EMC Corp	7.3.0.227	(M)
		SYMCFG	EMC Corp	7.3.0.227	(R)
		GNSdaemon	EMC Corp	7.3.0.227	(R)
		GNSdaemon	EMC Corp	7.3.0.228	(R)
		SYMCFG	EMC Corp	7.3.0.226	(R)
		GNSdaemon	EMC Corp	7.3.0.224	(R)
		SYMCFG	EMC Corp	7.3.0.224	(R)
		SYMCFG	EMC Corp	7.3.0.193	(R)
		SYMRDF	EMC Corp	7.3.0.189	(M)
		SYMCFG	EMC Corp	7.3.0.189	(R)

Legend for Attribute(s):
 (M) : Application Registered Remotely via RDF links (Multi-Hop).
 (R) : Application Registered Remotely via RDF links (One-Hop).

List host connections to arrays

Options

`-connections`

Returns the host connections to the array. Only hosts that have at least one registered application are listed.

Examples

To list the host connections to array 097 enter:

NOTE: To view the host connections for all arrays, omit the `-sid` option.

```
symcfg list -connections -capacity -sid 097
```

Sample output

NOTE: When displaying all arrays, the output is sorted according to each array

Host		Symmetrix		Capacity (GB)		
Node Name	ID	Director	Port	Mapped STDs	Mapped BCVs	Paired BCVs
6I	000190100097	FA-3B	0	135.0	67.4	8.4
6I totals:				130.8	54.8	8.4
LBQA0074	000190100097	FA-3B	0	135.0	67.4	8.4
LBQA0074 totals:				0.0	0.0	0.0
api105	000190100097	FA-3A	0	134.9	0.0	0.0
api105 totals:				134.9	0.0	0.0
api150	000190100097	FA-3A	0	134.9	0.0	0.0
api150 totals:				134.9	0.0	0.0
api31	000190100097	FA-7A	0	404.6	0.0	0.0
api31 totals:				404.6	0.0	0.0
lbqa0074	000190100097	FA-3B	0	135.0	67.4	8.4
lbqa0074 totals:				0.0	0.0	0.0

List PowerPath host registration records

Options

`-ppreg`

Returns the PowerPath host registration records stored on the target array.

Examples

To list the PowerPath host registration records on all arrays enter:

NOTE: To view the host connections for all arrays, omit the `-sid` option.

```
symcfg list -ppreg
```

Sample output

NOTE: When displaying all arrays, the output is sorted according to each array

```
Symmetrix ID      : 000197800188

Host Name:        hostabc
OS Version:       osver123
OS Revision:      osrev123
Hardware Vendor Name: hvn123
PowerPath Version: ppver123
PowerPath Patch Level: pppl123
PowerPath License Info: ppl123
Host Registration Time: 04/17/1979 00:16:15
Host Connectivity type: FC
Cluster Info:
  Cluster Name:    cluster123
  Cluster Node Name: clusnode123
WWNs:
  (1): 3132333435363738
VMs:
  (1)
```

```

VM Name      : vmname11
OS Vendor Info: vmosvendor11
(2)
VM Name      : vmname12
OS Vendor Info: vmosvendor12

Host Name:      hostxyz
OS Version:     osver567
OS Revision:    osrev567
Hardware Vendor Name: hvn567
PowerPath Version: ppver567
PowerPath Patch Level: pppl567
PowerPath License Info: ppl567
Host Registration Time: 07/19/2067 13:34:09
Host Connectivity type: FC
Cluster Info:
  Cluster Name:      cluster567
  Cluster Node Name: clusnode567
WWNs:
(1): 3132333435363738
VMs:
(1)
  VM Name      : vmname21
  OS Vendor Info: vmosvendor21
(2)
  VM Name      : vmname22
  OS Vendor Info: vmosvendor22
(3)
  VM Name      : vmname23
  OS Vendor Info: vmosvendor23
(4)
  VM Name      : vmname24
  OS Vendor Info: vmosvendor24
(5)
  VM Name      : vmname25
  OS Vendor Info: vmosvendor25
(6)
  VM Name      : vmname26
  OS Vendor Info: vmosvendor26
(7)
  VM Name      : vmname27
  OS Vendor Info: vmosvendor27
(8)
  VM Name      : vmname28
  OS Vendor Info: vmosvendor28
(9)
  VM Name      : vmname29
  OS Vendor Info: vmosvendor29
(10)
  VM Name      : vmname2a
  OS Vendor Info: vmosvendor2a

```

```
Symmetrix ID      : 000197100086
```

```
. . .
```

List host connections sorted by host names

Options

-ipv6

Displays a layout that does not truncate node names or addresses.

Examples

To list all of the host connections sorted by host names for array 097, enter:

```
symcfg list -connections -sorthost -sid 097
```

Sample output

Host		Symmetrix				
Node Name	IP Address	OS Name	OS Revision	ID	Director	Port
6I	172.23.195.65	WinNT	5.2.3790	000190100097	FA-3B	0
LBQA0074	172.23.195.74	WinNT	5.2.3790	000190100097	FA-3B	0
api105	172.23.192.105	SunOS	5.8	000190100097	FA-3A	0
api150	172.23.192.150	WinNT	5.2.3790	000190100097	FA-3A	0
api31	172.23.192.31	SunOS	5.8	000190100097	FA-7A	0

Supported director configuration types

Use the `symcfg list` command to gather information about the array directors. The following director types are supported:

- DA — Disk directors
- DX — External disk directors
- EA — ESCON directors
- SE — Gig-E directors
- EF — FICON (Fibre-ESCON) directors
- RA — SRDF directors
- RE — RDF Gig-E directors
- RF — RDF Fibre directors
- FA — Front-end (Fibre Channel) directors, including Fibre Channel over Ethernet (FCoE)
- IM — Infrastructure Manager (IM) directors
- EDS — Engenuity Data Services (EDS) directors

List director configuration data

Examples

To list configuration and status information about all directors on array 00019490004, enter:

```
symcfg list -dir ALL -sid 064
```

Sample output

```
Symmetrix ID: 000197100064 (Local)
```

S Y M M E T R I X		D I R E C T O R S			
Ident	Type	Engine	Cores	Ports	Status
IM-7A	IM	1	1	0	Online
IM-8A	IM	1	1	0	Online
ED-7B	EDS	2	2	0	Online
ED-8B	EDS	2	2	0	Online
DF-7C	DISK	3	2	4	Online
DF-8C	DISK	3	2	4	Online
FA-7E	FibreChannel	4	1	2	Online
FA-8E	FibreChannel	4	1	3	Online
RF-7G	RDF_BI_DIR	5	1	1	Online
RF-8G	RDF_BI_DIR	5	1	1	Online

List port flag information

Syntax

To list port flag information, use the following syntax:

```
symcfg [-sid <SymmID>][-offline]
list -port -free
<[-slot <#>] [-dx | -fa | -fcoe | -re | -rf | -se]
[-speed <#>] |
-dir <#>>
```

Options

- free** Lists the ports that are not associated with a director emulation.
- dx** Lists only DX (external) directors.
- fa** Lists only FA (Fibre) directors.
- fcoe** Lists only FCOE (Fibre Channel Over Ethernet) directors.
- re** Lists only RE (RDF Gig-E) directors.
- rf** Lists only RF (RDF Fibre) directors.
- se** Lists only SE (Gig-E) directors.

Support for multiple cores and ports

All director types can support multiple cores. The actual number of cores assigned to a director are statically configured. Use the `symcfg list - dir` command to view the number of cores assigned to a given director. In addition, all director types are capable of supporting a variable number of ports. Solutions Enabler supports reporting existing port associations, as well as the number and identity of ports available for association (free ports) and their respective supported interface types. A port's supported interface types determine which directors a given free port may be associated with.

List director information by type

Examples

To list all front-end directors (-FA) for array 005, enter:

```
symcfg list -FA ALL -sid 005
```

NOTE: Director types, -CA and -SA, are no longer supported. Special format displays such as those previously produced by the specification of -FA or -SE options are no longer supported. Two new director types are supported for Infrastructure Manager (IM) and Engenuity Data Services (EDS).

Sample output

```
Symmetrix ID: 000197300005 (Local)
S Y M M E T R I X      D I R E C T O R S
  Ident  Type          Engine  Cores  Ports  Status
-----
FA-1E   FibreChannel    1       6     5     Online
FA-2E   FibreChannel    1       4     4     Online
```

FA-3E	FibreChannel	2	6	5	Online
FA-4E	FibreChannel	2	3	2	Online

List director details by name and type

Options

-v (verbose)

Lists details about a specific director and type.

Examples

To display information about the Gig-E director 5G on array 064, enter:

```
symcfg list -SE 5G -sid 064 -v
```

Sample output

```
Symmetrix ID: 000197100064 (Local)
Time Zone      : EDT
Product Model  : VMAX40K
Symmetrix ID   : 000195700064
. . . .
Director Identification: SE-5G
Director Type      : GigE
Director Status    : Online
Director Symbolic Number : 05G
Director Numeric Number : 101
Director Engine Number : 1
Director Slot Number : 5
Number of Director Cores : 2
Number of Director Ports : 1
Director Port: 16
  WWN Port Name      : 5000097300039118
  Director Port Status : Online
  Negotiated Speed (Gb/Second) : 1
  Director Port Speed (Gb/Second) : 4
```

List director port data

Options

-port

Lists ports that are online or offline on SA, FA, SE, or RA directors.

-p <#>

Restricts output to a specific port number.

Examples

To list port status for all SE directors on array 230, enter:

```
symcfg list -port -SE ALL -sid 230 -detail
```

Sample output

```
Symmetrix ID: 000197100230
  S Y M M E T R I X      D I R E C T O R      P O R T S
  Ident  Port  WWN              Type              Flags  Speed  Status
  -----  ---  -----  -----  -----  -----  -----
  SE-9G   0  50000972C011C918  GigE              X--    8  Online
  SE-10G  0  50000972C011C919  GigE              X--    8  Online
  SE-9H   0  50000972C011C91A  GigE              X--    8  Online
  SE-10H  0  50000972C011C91B  GigE              X--    8  Online
```


Legend:

Flags:

(A) CLX Enabled : X = True, . = False, - = N/A
 (C) HAP Enabled : X = True, . = False, - = N/A
 (R) ADIUS Enabled : X = True, . = False, - = N/A

List addresses of devices mapped to directors

Options

-address

Identifies the address information for devices accessible through specific directors from a host-based view of the storage environment.

Examples

To list the address information for all director types on array 064, enter:

```
symcfg list -dir ALL -sid 064 -address
```

To list address information for port 10 on director FA 2E, enter:

```
symcfg list -sid 075 -FA 2E -address -p 10
```

Sample output

For all director types:

```
Symmetrix ID: 000197100064 (Local)
  Director          Device Name          Attr          Address
-----
 Ident   Port  Sym   Physical          VBUS  TID  LUN
-----
 FA-8F   10  807A  Not Visible          0    00  000
          807B  Not Visible          0    00  001
          807C  Not Visible          0    00  002
          807D  Not Visible          0    00  003
```

For port 10 on director type FA:

```
Symmetrix ID: 000197100075 (Local)
  Director          Device Name          Attr          Address
-----
 Ident   Port  Sym   Physical          VBUS  TID  LUN
-----
 FA-2E   10  000AC  Not Visible    ACLX    0    00  100
          00040  Not Visible    N/A    N/A  N/A
          00041  Not Visible    N/A    N/A  N/A
          00042  Not Visible    N/A    N/A  N/A

Total
Mapped Devices:          4
Including Metamembers:   4
Port Available Addresses: 4092
Director Available Addresses: 98284 (s)
Legend for Available address:
(s): The Available Addresses for a director are shared among
its ports (shared)
```

NOTE: When LUN information is not available, VBUS TID, and LUN address are reported as N/A.

List the next available device address

Options

-available

Returns the next available address that can be used for a device.

Examples

To list next available address on array 097, enter:

```
symcfg list -dir all -address -available -sid 097
```

Sample output

NOTE: VBUS, TID, and LUN address values with an asterisk (*) represent a gap in the address assignments, or are the next available address in the run.

```
Symmetrix ID: 000190100097
Director
-----
Ident      Symbolic  Port  Sym   Physical          Attr   Address
-----
FA-3A     03A      0     0E2A  /dev/vx/rdmp/c5t0d0s2  0     00  000
          0E2B  /dev/vx/rdmp/c5t0d1s2  0     00  001
          04AA  /dev/vx/rdmp/c5t0d2s2  (M)   0     00  002
          04AE  /dev/vx/rdmp/c5t0d3s2  (M)   0     00  003
          04B2  /dev/vx/rdmp/c5t0d4s2  (M)   0     00  004
          04B6  /dev/vx/rdmp/c5t0d5s2  (M)   0     00  005
          04BA  /dev/vx/rdmp/c5t0d6s2  (M)   0     00  006
          04BE  /dev/vx/rdmp/c5t0d7s2  (M)   0     00  007
          04C2  /dev/vx/rdmp/c5t0d8s2  (M)   0     00  008
          04C6  /dev/vx/rdmp/c5t0d9s2  (M)   0     00  009
          -    AVAILABLE             0     00  00A *
```

Take RA directors offline

Examples

To take RA-12 for array 097 offline, enter:

```
symcfg offline -RA 12 -sid 097
```

Bring RA directors online

Examples

To bring RA director 12 online for array 097, enter:

```
symcfg online -RA 12 -sid 097
```

Take front-end director ports offline

Examples

To take port 4 of SA-12 in array 097 offline, enter:

```
symcfg offline -SA 12 -P 4 -sid 097
```

NOTE: Do not turn off the only connection from the host to the array, otherwise another host that has connection to the array must be used to bring a director port back online.

Bring front-end director ports online

Examples

To bring port 4 of SA-12 in array 097 back online, enter:

```
symcfg online -SA 12 -P 4 -sid 097
```

Array external locks

Array external locks are used by SYMAPI (locks 0 to 15) and also for applications assigned by EMC (>15) to lock access to the entire array during critical operations. (Base SRDF operations use lock 0 and the Optimizer uses lock 13.) Use the `symcfg list -lockn` command, to list all locks on one or all arrays or just the locks targeted to specific operations.

List all array locks

Examples

To return a list of all host-visible arrays (local and remote), along with details about all array exclusive locks, enter:

```
symcfg list -lockn all
```

Sample Output

The returned list contains three local arrays that have no known locks, as specified by the N/A values. Remote array 000187900039 has an exclusive lock number 15 for a configuration change activity (ConfigChg).

S Y M M E T R I X		L O C K S			
SymmID	Attachment	Lock Status	Lock Number	Lock Usage	Time Held (Sec)
000000006196	Local	Unknown	N/A	N/A	N/A
000184600063	Local	Unknown	N/A	N/A	N/A
000184600282	Local	Unknown	N/A	N/A	N/A
000187900039	Remote	EXCLUSIVE	15	ConfigChg	307

In the previous example,

List lock number details

Examples

To return a list of all host-visible arrays (local and remote), and details about lock 0, enter:

```
symcfg list -lockn 0
```

Sample output

NOTE: If an array is holding a lock other than 0, the output still returns a lock number of N/A.

S Y M M E T R I X		L O C K S			
SymmID	Attachment	Lock Status	Lock Number	Lock Usage	Time Held (Sec)
000000006196	Local	Unknown	N/A	N/A	N/A
000184600063	Local	Unknown	N/A	N/A	N/A
000184600282	Local	Unknown	N/A	N/A	N/A
000187900039	Remote	Unknown	N/A	N/A	N/A

List lock details

Options

-v	Lists extended lock information about the application and host that owns the lock
all	Lists locks for all arrays.
#	Lists specific lock number.
RDF	Lists SRDF locks only.
RDFA	Lists SRDF/A locks only.
SRDF_MSCS	Lists SRDF MSCS locks only.
GNS	Lists GNS locks only.

Examples

To list information about lock 23, enter:

```
symcfg list -lockn 23 -sid 190
```

To return verbose information about lock 15, enter:

```
symcfg list -lockn 15 -sid 207 -v
```

Sample output

Output for information on lock 23.

```
Symmetrix ID: 000000006190
      S Y M M E T R I X      L O C K S
      Lock      Lock      Lock      Time
      SymmID      Attachment      Status      Number      Usage      Held (Sec)
000000006190 Local      Locked      23      Unknown      20
```

Output for verbose information on lock 15.

```
Symmetrix ID: 000192600207
Symmetrix ID      : 000192600207 (local)
Lock Number      : 15
Lock Usage      : Config Change
Time Held in Seconds : 1801
Lock Holder ID   : 0xf000128e
Initiator        : 0
Director Number  : 55
Logical Path ID  : -1
Lock Owner       : Config Server
Config Change Session ID : 4356
Migration Session Name : N/A
```

Release an array lock

Description

The `symcfg release` command determines the lock owner and releases the lock. If held by an SRDF control operation, or by the configuration change server, SYMCLI blocks the user command and displays a failure message identifying the lock owner. The lock owner

can be a configuration change session, a migrate session, or an internal change. Using the `symconfigure` or the `symmigrate` command, abort the correct session that owns the lock.

NOTE: Releasing a lock on an array is available but not recommended unless it is confirmed that the lock is stranded.

For information about device external locks that only target specific devices, refer to [Device external locks](#) on page 358.

Examples

To release an external lock 0, which has confirmed to be stranded on array 097, enter:

```
symcfg -sid 097 -lockn 0 release
```

List all LRU cache management groups

Examples

To view a list of all LRUs for array 6196, enter:

```
symcfg list -lru all -sid 6196
```

Sample output

```
Symmetrix ID: 000000006196 (Local)
```

LRU Num	LRU Name	Cache Slots	Percent of Total
1	GROUP_0	11971	4%
2	GROUP_1	11971	4%
3	GROUP_2	11971	4%
4	GROUP_3	11971	4%
5	GROUP_4	11971	4%
6	GROUP_5	11971	4%
7	GROUP_6	11971	4%
8	GROUP_7	11971	4%
9	GROUP_8	11971	4%
10	GROUP_9	11971	4%
11	GROUP_A	11971	4%
12	GROUP_B	11971	4%
13	GROUP_C	11971	4%
14	GROUP_D	11971	4%
15	GROUP_E	11971	4%
16	GROUP_F	11971	4%
Total		278616	

Network configuration file

The Solutions Enabler configuration file `netcnfg` is used by the SYMCLI client library to specify attributes of a remote server where array management operations should be directed. Each line in the file associates a user-chosen service name with host name or IP address, port, and security level information.

The `netcnfg` file is a template and an editable file located in the SYMAPI configuration directory. The location of this directory varies according to operating system. The *Solutions Enabler Installation Guide* explains how to set up a Solutions Enabler client and configure services in the `netcnfg` file.

List SYMAPI services

Description

The `symcfg -services list` command validates the syntax of the `netcnfg` file entries and displays the configured network services available for use by the SYMAPI client connection.

Examples

```
symcfg list -services
```

Sample output

```
SYMAPI    NET    SERVICES

Name      Pairing
          Method  Type    Node Name      Address      Port  Security
-----
single_entry Single  TCPIP   vmax1.mmcs1.xyz.com 2707 SECURE
ord_guest  Ordered TCPIP   vmax1.mmcs1.com    2707 SECURE
ord_guest  Ordered TCPIP   vmax1.mmcs2.com    2707 SECURE
bal_guest  Balanced TCPIP   vmax1.mmcs1.xyz.com 2707 SECURE
bal_guest  Balanced TCPIP   vmax1.mmcs2.com    2707 SECURE
```

Pairing Method descriptions:

- Single — Single entry service name.
- Ordered — The SYMAPI client library first attempts a client/server session with the server named as the first of the two entries in the `netcnfg` file. If that attempt fails, the library tries the second entry.
- Balanced — The SYMAPI client library applies a random number to select the first entry to attempt a client/server session. So it may select the first or the second entry in the `netcnfg` file.

List memory board information

Description

Memory board information includes the number of boards, the slot number, and the capacity information in MBs.

Examples

To view the available memory board information for all arrays, enter:

```
symcfg list -memory
```

Mainframe CU image and split information reporting

The `symcfg list` command is used to report mainframe CU (Controller Unit) image and split information if the storage environment contains devices mapped to either EA (ESCON) or EF (FICON) front-end directors. Since devices in the mainframe environment are managed with respect to the CU image that they are a part of, SYMCLI creates a view of the CU images that are defined within the array. A CU image definition includes the SSID assigned to the image, the split name, the front-end ports to which it is mapped, the devices included in the image, and their base and alias addresses. It also indicates whether it uses dynamic or static PAV (parallel access volumes) and whether the CU is online or not.

The `ficon_split` option reports split information on the array.

List mainframe CU images

Examples

To list CU images for array 086, enter:

```
symcfg -sid 086 list -cuimage
```

Sample output

```
Symmetrix ID      : 000197100086
PAV Aliasing      : DynamicStandardPAV
```

```

CU image number: 0x00

Sub System ID           : 0x0140
Split Name              : split0

CU status               : N/A
PAV Aliasing            : DynamicStandardPAV
Director Port Assignments (02) : EF-01H:24
                        : EF-01H:26

Number of Devices       : 0
Device Ranges           : N/A
Number of Base Addresses : 0
Number of Alias Addresses : 0

```

```

CU image number: 0x01

Sub System ID           : 0x0141
Split Name              : split0

CU status               : N/A
PAV Aliasing            : DynamicStandardPAV
Director Port Assignments (02) : EF-01H:24
                        : EF-01H:26

Number of Devices       : 10
Device Ranges           : 0x10 - 0x19
Number of Base Addresses : 10
Number of Alias Addresses : 0

```

Show mainframe CU images

Examples

To show CU images for array 086, enter:

```
symcfg -sid 086 show -cuimage 1
```

Sample output

```

Symmetrix ID           : 000197100086

CU image number       : 0x01
Sub System ID         : 0x0060
Split Name            : split0

CU status              : Offline
PAV Aliasing          : DynamicStandardPAV
Director Port Assignments (02) : EF-01H:24
                        : EF-01H:26

Number of Devices      : 2
Number of Base Addresses : 2
Number of Aliases Addresses : 0
Range of Alias Addresses : 31 - 30

SymDev      Base Address
-----
0032A      120
00D24      130

CU image number       : 0x01
Sub System ID         : 0x0141
Split Name            : split1

CU status              : Offline
PAV Aliasing          : DynamicStandardPAV
Director Port Assignments (02) : EF-01H:25
                        : EF-01H:27

```

```

Number of Devices      : 3
Number of Base Addresses : 3
Number of Aliases Addresses : 0
Range of Alias Addresses : 41 - 50

```

```

SymDev      Base Address
-----
00D16      120
00D17      130
00D18      140

```

List mainframe splits

Syntax

```

symcfg [-sid <SymmID>] [-offline]
list -ficon_split [-v]
show -ficon_split <Split Name>
list -ficon_split -address [-available]

```

Options

- ficon_split**
Lists splits.
- v (-verbose)**
Lists split details (ie: CU image number, Ficon port assignments)
- address (-avail)**
Lists base address for split.
- available (-addr)**
Lists device base address available for split.

Examples

To list splits for array 086, enter:

```
symcfg sid 086 list -ficon_split
```

To list splits, device base address, and base address for devices available for split, for array 086, enter:

```
symcfg sid 086 list -ficon_split -addr -avail
```

Sample output

Output for splits:

```

S P L I T S

Symmetrix ID: 000197100086
-----
Split Name      Serial Flg
                #      M
-----
split0          123456 H
split1          123457 H
split2          123458 D

Legend:
Flags:
(M)ode H = Hyper PAV, D = Dynamic PAV

```

Output for split device base address, and base address for devices available for split:

```

Split      CU      Device Name      Base
-----

```


Name	Img #	Sym	Physical	Addr
split0	01	00015	/dev/sdj	000
		00016	/dev/sdk	001
		00017	/dev/sdl	002
		-	AVAILABLE	003
		-	AVAILABLE	004

Total
Mapped Devices: 3
CU Available Addresses: 20
Split Available Addresses: 6344 (s)

Split	CU	Device Name	Base
-------	----	-------------	------

Name	Img #	Sym	Physical	Addr
split0	02	00018	/dev/sdm	010
		00019	/dev/sdn	011
		-	AVAILABLE	012

Total
Mapped Devices: 2
CU Available Addresses: 20
Split Available Addresses: 6344 (s)

Split	CU	Device Name	Base
-------	----	-------------	------

Name	Img #	Sym	Physical	Addr
split1	03	00020	/dev/sdy	020
		00021	/dev/sdz	021
		-	AVAILABLE	022

Total
Mapped Devices: 2
CU Available Addresses: 20
Split Available Addresses: 6344 (s)

Legend for Available address:

(s): The Available Addresses for a split are shared among its ports (shared)

List operating system patches

Examples

To list all of the OS patches installed on a array 207, enter:

```
symcfg list -upatches -sid 207
```

Environmental data

Use the `list -env_data` option to list the status of the major hardware modules including fans, power supplies, drive enclosures, and link control cards. An array ID must be specified when querying for environmental data.

List all environment data on array

Examples

To list an overall status for all environmental components on array 150, enter:

```
symcfg -sid 150 list -env_data
```

Show environmental data on array component

Examples

To return a detailed status for each environmental component for SystemBay on array 150, enter:

```
symcfg -sid 150 show -env_data SystemBay
```

List environmental data for specific service state

Options

`-service_state`

Returns data for possible service states: degraded, failed, or normal. To list all service states except one prefix the service state value with not, such as `-service_state notfailed`

Examples

To list the environmental data for array 150 with a service state of failed, enter:

```
symcfg -sid 150 list -env_data -service_state failed
```

 **NOTE:** Returned data only contains information about the bay containing the failure, and the components in the failed state.

Listing array environmental data example

Examples

To list environmental data details for array 64, enter:

```
symcfg list -env_data -v -sid 64
```

Sample output

```
Symmetrix ID           : 000195700064
Timestamp of Status Data : 09/21/2011 14:01:51
System Bays
Bay Name                : SB-1
  Bay LED state          : Normal (On)
  Front Door Bay LED state : Normal (On)
  Rear Door Bay LED state : Normal
  Number of Standby Power Supplies : 2
  Number of Drive Enclosures : 1
  Number of Enclosure Slots : 1
  Number of MIBE Enclosures : 2
  Status of Contained Modules
  Standby Power Supplies
    SPS-1A (Aggregate) : Normal
    SPS-TRAY-1A        : Normal
    SPS-BATTERY-1A     : Normal
  SPS-1B (Aggregate)   : Normal
    SPS-TRAY-1B        : Normal
    SPS-BATTERY-1B     : Normal
  Enclosure Slot Number : 1
  Enclosure Slot State  : Normal
  MM-7                  : Normal
  MM-8                  : Normal
  DIR-1                 : Normal
    PS-A                : Normal
    FAN-1               : Normal
    BOOT-DRIVE-0        : Normal
  DIR-2                 : Normal
    PS-B                : Normal
    FAN-2               : Normal
```

```

        BOOT-DRIVE-0                : Normal
    Drive Enclosure Number          : 1
        Drive Enclosure State       : Normal
        SSC                          : Normal
        LCC-A                        : Normal
        LCC-B                        : Normal
        ICM-A                        : Normal
        ICM-B                        : Normal
        PS-A                         : Normal
        PS-B                         : Normal
        FAN-1                        : Normal
        FAN-2                        : Normal
    MIBE Name                       : MIBE-A
        MIBE State                   : Normal
        PS-A                         : Normal
        PS-B                         : Normal
        CM                           : Normal
    MIBE Name                       : MIBE-B
        MIBE State                   : Normal
        PS-A                         : Normal
        PS-B                         : Normal
        CM                           : Normal
Drive Bays
    Bay Name                       : DB-1A
    Bay LED state                   : Normal (On)
    Number of Standby Power Supplies : 4
    Number of Drive Enclosures      : 16
    Status of Contained Modules
        Standby Power Supplies
            SPS-1A (Aggregate)       : Normal
            SPS-TRAY-1A             : Normal
            SPS-BATTERY-1A          : Normal
            SPS-1B                   : Normal
            SPS-4A                   : Normal
            SPS-4B                   : Normal
        Enclosure Number             : 2
        Enclosure State             : Normal
        MM-A                        : Normal
        MM-B                        : Normal
        DIR-3                       : Normal
DIR-4                               : Normal

```

List device pools

Syntax

```

symcfg [-sid SymmID] [-offline] [-mb | -gb | -tb]
        [-i Interval] [-c Count]
list [-pool [-snap][-rdfa_dse [-rdfg GrpNum]][-thin]
        [-fba] [-ckd] [-ckd3390] [-ckd3380] [-as400] [-all] [-v]]

```

Options

-i *Interval* -c *Count*

Checks status of the pool(s) continuously for a certain period of time.

-v

List details of each pool in the returned data set. It is equivalent to using the *GrpNum* command on all desired pools.

-mb | -gb | -tb

By default, the space consumption of devices and pools is shown as a number of tracks. For the `symcfg list` and `symcfg show` commands, the output is shown in megabytes, gigabytes, or terabytes by specifying one of these options. The gigabytes display has one decimal point precision.

- pool**
Lists information for TF/Snap, SRDFA/DSE, and thin pools in a common output.
- snap**
Lists information for TF/Snap pools only.
- rdfa_dse**
Lists information for SRDF/A DSE pools only.
- rdfg**
Used with `-rdfa_dse GrpNum` to limit the display to the SRDF/A DSE pools that are related to the specified SRDF group. This includes pools that are associated with the group and pools that have been disassociated from the group, but may still have some data for the group.
- thin**
Displays information about thin device pools only.
- all**
Includes both enabled and disabled devices in the calculations for `Free Tracks` and `Full %`. Otherwise, only enabled devices are included. When specifying the `-all` option, both enabled and disabled devices are included in the calculations of `Free Tracks` and `Full %`. Otherwise, only enabled devices are included. For example, without the `-all` option, the `Free Tracks` field include free tracks from all enabled devices and the `Full %` field would be based on the Enabled tracks. With the `-all` option specified, the `Free Tracks` field include free tracks from both enabled and disabled devices and the `Full %` would be based on the Usable Tracks.
- fba | -ckd | -ckd3390 | -ckd3380 | -as400**
Filters the pool display to the specified emulation type.

Examples

To list details about all thin pools in array 087:

```
symcfg list -pool -sid 087 -thin -all -detail
```

Sample output

```
Symmetrix ID: 000197100087
SYMMETRIX THIN POOLS
-----
Pool      Flags  Dev      Total    Usable    Free      Used Full Subs Comp   Shared
Name      PTECSL Config  Tracks  Tracks   Tracks   Tracks  (%) (%) (%)   Tracks
-----
DG1_CKD10K TFC-EI 2-Way Mir 6274800 6274800 6274800    0  0  0  0    0
DG1_FBA10K TFF-EI 2-Way Mir 11536560 11536560 2469946 9066614 78  0  0    0
DG2_FBA7_2 TSF-EI RAID-5(3+1) 21546000 21546000 21546000    0  0  0  0    0
DG3_FBA7_2 TSF-EI 2-Way Mir 3591000 3591000 3591000    0  0  0  0    0
Total
Tracks      42948360 42948360 33881746 9066614 21  0  0    0
Legend:
(P)ool Type:
  S = Snap, R = Rdfa DSE T = Thin
(T)echnology:
  S = SATA, F = Fibre Channel, E = Enterprise Flash Drive, M = Mixed, - = N/A
Dev (E)mulation:
  F = FBA, A = AS400, 8 = CKD3380, 9 = CKD3390, - = N/A
(C)ompression:
  E = Enabled, D = Disabled, N = Enabling, S = Disabling, - = N/A
(S)tate:
  E = Enabled, D = Disabled, B = Balancing
Disk (L)ocation:
  I = Internal, X = External, M = Mixed, - = N/A
```

 **NOTE:** The pool state indicates whether there are any devices enabled in the pool.

View Snapshot policy details

Syntax

```
symcfg [-sid <SymmID>]
list -policy -type <snap> [-detail]
```

```
show -policy <PolicyName> -type <snap>
```

Options

-type

Specifies the type of Snapshot policy. <snap> is the only option supported at the current Solutions Enabler version.

-policy <PolicyName>

Specifies the Snapshot policy you want to list/show.

Examples

To list all Snapshot policies that are defined on array 084, enter:

```
symcfg list -policy -type snap -sid 084
```

```
Symmetrix ID          : 000197100084

Snapshot Policy Name   Flags  Interval  Offset  Max
-----
FINANCE_Daily_4PM     XX.    1 0 0    0 4 0    100  Every day at 4:00
FINANCE_Daily_10AM    X..    1 0 0    0 10 0   100  Every day at 10:00
FINANCE_HOURLY        ...    0 1 0    0 0 0    50   Every hour, starting at 0:00
FINANCE_WEEKLY        X..    7 0 0    2 0 0    10   Every Wednesday at 0:00
hr_sg_Daily_10AM      X..    1 0 0    0 0 0    200  Every day at 0:00  payroll_sg_Daily_2AM  XX.  1 0 0 0 0 2
0_1000_Every day at 2:00
payroll_sg_Daily_5PM  X..    1 0 0    0 17 0   200  Every day at 17:00

Legend:
Flags:
(S)ecure           X = Secured, . = Non-Secured
Sus(P)ended        X = Suspended, . = Not Suspended
(C)loud            X = Yes, . = No, - = N/A
```

To list all Snapshot Protection Policies defined on the array in detail, enter:

```
symcfg list -policy -type snap -v -sid 084
```

```
Symmetrix ID          : 000197100084

Snapshot Policy Name   : FINANCE_Daily_4AM
  Interval (D:HH:MM)   : 1:00:00
  Base Offset (D:HH:MM) : 0:04:00
  Maximum Count        : 100
  Secure Setting       : Yes
  Suspended            : No
  Critical Compliance  : N/A
  Warning Compliance   : N/A
  Cloud Policy         : No
  Last Update Time     : Wed Jan 12 07:52:02 2018
  Last Used Time       : Wed Mar 12 04:01:00 2018
  Schedule Description : Every day at 4:00

Associated SGs (1)
{
-----
Storage Group Name      Storage Group
-----
finance_sg_child2      Suspension
-----
                          Yes
}
. . .

Snapshot Policy Name   : FINANCE_Daily_10AM
  Interval (D:HH:MM)   : 1:00:00
  Base Offset (D:HH:MM) : 0:10:00
  Maximum Count        : 100
  Secure Setting       : Yes
  Suspended            : Yes
```

```

Critical Compliance      : N/A
Warning Compliance      : N/A
Cloud Policy            : No
Last Update Time        : Wed Jan 12 07:52:02 2018
Last Used Time          : Wed Mar 12 10:01:00 2018
Schedule Description    : Every day at 10:00

```

```
Associated SGs (3)
```

```

{
-----
Storage Group Name      Storage Group
                        Suspension
-----
finance_sg              No
finance_sg_child1      No
finance_sg_child2      No
}

```

```

. . .
Snapshot Policy Name    : FINANCE_WEEKLY
Interval (D:HH:MM)     : 7:00:00
Base Offset (D:HH:MM)  : 02:00:00
Maximum Count          : 10
Secure Setting         : Yes
Suspended              : No
Critical Compliance     : N/A
Warning Compliance     : N/A
Cloud Policy           : No
Last Update Time       : Wed Jan 15 09:51:04 2018
Last Used Time         : Wed Jan 17 10:31:00 2018
Schedule Description    : Every Wednesday at 0:00

```

```
Associated SGs (3)
```

```

{
-----
Storage Group Name      Storage Group
                        Suspension
-----
finance_sg              No
finance_sg_child1      No
finance_sg_child2      Yes
}
. . .

```

To show details of a specified Snapshot Protection Policy defined on array 084, enter:

```
symcfg show -policy finance_daily_4pm -type snap -sid 084
```

```

Snapshot Policy Name    : FINANCE_Daily_4AM
Interval (D:HH:MM)     : 1:00:00
Base Offset (D:HH:MM)  : 0:04:00
Maximum Count          : 100
Secure Setting         : Yes
Suspended              : No
Critical Compliance     : N/A
Warning Compliance     : N/A
Cloud Policy           : No
Last Update Time       : Wed Jan 12 07:52:02 2018
Last Used Time         : Wed Mar 12 04:01:00 2018
Schedule Description    : Every day at 4:00

```

```
Associated SGs (1)
```

```

{
-----
Storage Group Name      Storage Group
                        Suspension
-----
finance_sg_child2      Yes
}

```

Show thin pool rebalancing

Syntax

To show pool rebalancing parameters, use the following syntax:

```
symcfg show -thin -pool PoolName -detail -all -sid SymmID
```

NOTE: Pool rebalancing is for thin pools only and is not applicable to Snap or SRDF/A DSE pools.

Examples

To show pool rebalancing parameters for thin pool Mig_trg2 on array 432, enter:

```
symcfg show -pool Mig_trg2 -thin -sid 432 -all -detail
```

Sample output

```
Symmetrix ID: 000194900432
Symmetrix ID           : 000194900432
Pool Name              : Mig_trg2
Pool Type              : Thin
Disk Location          : External
Technology             : N/A
Dev Emulation          : FBA
Dev Configuration     : 2-Way Mir
Pool State             : Enabled
Compression State     : Enabled
# of Devices in Pool  : 10
# of Enabled Devices in Pool : 10
# of Usable Tracks in Pool : 99000
# of Allocated Tracks in Pool : 14081
# of Thin Device Tracks : 11549
# of DSE Tracks       : 2298
# of Local Replication Tracks : 234
# of Tracks saved by compression : 0
# of Shared Tracks in Pool : 0
Pool Utilization (%)  : 3
Pool Compression Ratio (%) : 0
Max. Subscription Percent : N/A
Rebalance Variance   : N/A
Max devs per rebalance scan : N/A
Pool Reserved Capacity : N/A
. . .
Legend:
Enabled devices FLG:
(S)hared Tracks : X = Shared Tracks , . = No Shared Tracks
Bound devices FLG:
S(T)atus       : B = Bound, I = Binding, U = Unbinding, A = Allocating,
                 D = Deallocating, R = Reclaiming, C = Compressing,
                 N = Uncompressing, . = Unbound
```

NOTE: The FLG S flag field indicates whether or not there are shared allocations on each device as noted in the above legend.

Table 22. Pool Rebalancing Parameters

Pool Rebalancing Parameters	Description
Rebalancing variance	Targets the device utilization variance for the rebalancing algorithm. The rebalancing algorithm attempts to level distribution of data in a pool so that the percentage utilization of any device in the pool is within the target variance of the percentage utilization of any other device in the pool.

Table 22. Pool Rebalancing Parameters (continued)

Pool Rebalancing Parameters	Description
Max devs per rebalance scan	Lists the maximum number of devices in a pool to use in the rebalancing algorithm. The default is 256.

List feature registrations and usage data

Description

The `symcfg list -feature` command lists feature class registrations and usage data for a specified array. Where appropriate, capacity types and limits are also displayed.

Options

- `-v`
Displays usage information.
- `-class`
Limits the feature display to a specified class of features.

Examples

To list feature class registrations for Local Replication class on array 341, enter:

```
symcfg list -feature -sid 341 -class Local Replication
```

Sample output

```
Symmetrix ID : 000194900341
  Feature Name      : TimeFinder/Mirror
  Feature Type      : Product
  Feature Class     : Local Replication
  Feature Capacity Type : TB of Total Capacity
  Feature Capacity  : 0
  SATA Capacity    : 0
  Enabled Status    : Disabled
  Enabled Change Date : 05-Jan-2011 17:10
  Feature Name      : SYMM_TF_CLONE
  Feature Type      : Product
  Feature Class     : Local Replication
  Feature Capacity Type : TB of Registered Capacity
  Feature Capacity  : 500
  SATA Capacity    : 0
  Enabled Status    : Disabled
  Enabled Change Date : 10-Mar-2011 14:32
  Feature Name      : SYMM_TF_SNAP
  Feature Type      : Product
  Feature Class     : Local Replication
  Feature Capacity Type : TB of Total Capacity
  Feature Capacity  : 100
  SATA Capacity    : 500
  Enabled Status    : Disabled
  Enabled Change Date : 10-Mar-2011 14:30
```

Device-level data

From the perspective of software running on a host system, an array is many *physical devices* connected to one or more I/O controllers. A host application addresses each of these devices using a physical device name. Each physical device defined in the configuration database has a specific set of attributes (such as vendor ID, product ID, revision level, and serial ID).

A device can map to a part of a physical disk or to an entire disk. The part of a physical disk to which a device is mapped is called a *hypervolume* or a *hyper*. A device may map to multiple hypers (containing identical copies of data) depending on its mirror configuration.

The array database file maintains device-level configuration and status information for each device on every array that is accessible from the host. Using SYMCLI, a list of all available devices can be displayed. The listed device data is used to obtain configuration and status information. This information identifies back-end information for the device's disk directors and corresponding hypervolumes, and their mappings to disk drives.

Device types

SYMCLI defines and configures devices for numerous specialized roles defined as *device types*. Each device type has specialized characteristics that enables a device to participate in various SYMAPI operations.

Device type	Description
ACLX devices	<p>ACLX devices are device masking devices similar to VCM devices that are used for storage provisioning using auto-provisioning groups.</p> <p>VMAX3 arrays come pre-configured with one ACLX Device. The ACLX attribute will cannot be removed, and the ACLX devices cannot be created or deleted. The device will be visible to hosts at the LUN Address configured in the ACLX Device Lun address global configuration, and zero is the default value. The ACLX device is only visible on front end ports where the <code>Show_ACLX_Device</code> port attribute is enabled.</p>
Data Domain devices	An encapsulated Data Domain device.
Gatekeeper devices	Gatekeeper devices are LUNs that act as the target of command requests to microcode-based functionality. For detailed information on gatekeeper management, refer to the <i>EMC Solutions Enabler Installation Guide</i> .
RDF devices	Devices configured as RDF1, RDF2, or RDF21 to support SRDF operations. SRDF is a business continuance solution that maintains a device-level mirror of array data on remotely attached arrays. These arrays also may be located in physically separate sites. SRDF provides a recovery solution for component or site failures using remotely mirrored devices. SRDF reduces backup and recovery costs and significantly reduces recovery time after a disaster. For more information, refer to the <i>EMC Solutions Enabler SRDF CLI Product Guide</i> .
Thin devices (TDEV)	Thin devices used for Thin Provisioning are devices that do not have storage allocated to them when they are created. To a host operating system, they look like regular devices with their configured capacity. The host treats them as regular devices and writes and reads from these devices like regular devices.
RecoverPoint Splitter devices	Devices that are in use by the RecoverPoint Splitter. With HYPERMAX OS 5977, indicates the specific RecoverPoint device type as production, replica or internal.

List array device information

Description

Use the following two commands to list information about array devices:

- `sympd` — Lists array devices that are host-visible.

- `symdev` — Displays information about all array devices, host-visible or not.

Syntax

To list host-visible devices, use the following syntax:

```
symprd [-offline] [-sid <SymmID>] [-v]
  list [-resv]
  list [-SA <#|ALL>] [-p <#>] [-scsi] [-fibre]
    [-escon] [-ficon] [-gige] [-iscsi_port <#>]
    [-powerpath] [-vcm | -aclx] [-pdevfile] [-cyl]

symprd [-offline] [-sid <SymmID>] [-v]
  list [-DA <#|ALL>] [-interface <#|ALL>]
    [-disk <#|ALL>] [-hyper <#|ALL>]
    [-spindle]
  list [-DX <#|ALL>] [-spindle]
  list [-vm]
```

Options

- sid** Limit the device output to a specific array.
- DA** List host-visible array devices that match a specific DA.
- DX** Lists host-visible array devices that match a specific DX.
- interface** Lists host-visible array devices that match a specific interface.
- disk** Lists host-visible array devices that match a specific disk.
- hyper** Lists host-visible array devices that match a specific hyper-volume.
- spindle** Includes Spindle ID Information.
- vm** Displays valid virtual machine names on VMware ESX server environments.
- SA** List the host-visible array devices that match a specific front-end director number.
- P #** List the host-visible array devices that match a specific front-end director port number.
- scsi** List the host-visible array devices that are mapped to SCSI front-end directors.
- fibre** List the host-visible array devices that are mapped to Fibre front-end directors.
- ficon** List the host-visible array devices that are mapped to FICON front-end directors.
- escon** List the host-visible array devices that are mapped to ESCON front-end directors.
- gige** List the host-visible array devices that are mapped to Gig-E front-end directors.

NOTE: The `symprd` command only provides a limited amount of filter options. Use the `symdev list pd` command for additional filter options on physical devices. Refer to [Filter list for device data](#) on page 355 for more information.

Examples

To list host-visible array devices that match a front-end director 07E for array 064, enter:

```
sympd list -SA 7E -sid 064
```

Sample output

Array device names are listed in the `Sym` column.

```
Symmetrix ID: 000197100064
      Device Name           Dir           Device
-----
Physical           Sym   SA :P   Config           Attribute   Sts   Cap
-----
/dev/sdt           08028 07E:008 TDEV           N/Grp'd    RW   2063
/dev/sdu           08029 07E:008 TDEV           N/Grp'd    RW   2063
/dev/sdv           0802A 07E:008 TDEV           N/Grp'd    RW   2063
```

Report spindle ID information

The following commands report spindle ID information:

- `sympd show`
- `sympd list -v`
- `symdev list`
- `symdev show`

sympd show

- For all but unprotected devices, the spindle ID reports as part of the RAID Group Information section.
- For unprotected devices, the spindle ID reports in both the RAID Group Information section and the Back-end Disk Director Information section.
- For non-RAID devices (VAULT devices), spindle IDs report in the Back-end Disk Director Information section.

sympd list -v

Reports spindle ID information as part of both the back-end disk director Information section and the RAID group Information section.

symdev list

- Lists host-visible array devices for the following specified values:
 - DA (`-da`)
 - DX (`-dx`)
 - interface (`-interface`)
 - disk (`-disk`)
 - hyper-volume (`-hyper`)
- When specified along with the `-spindle` option, the reports Spindle ID information.

symdev show

- For all but unprotected devices, the spindle ID reports as part of the RAID Group Information section.
- For unprotected devices, the spindle ID reports in both the RAID Group Information and the Back-end Disk Director Information section.
- For non-RAID devices (VAULT devices), spindle IDs reports in the Back-end Disk Director Information section.
- Provides similar output as the `sympd` command, but includes all array devices and lists them by array device names.

External spindle devices

FAST.X allows for an external disk to attach external storage to VMAX Family arrays. Adding an eDisk to an array makes the eDisk's capacity available to the array as an external spindle. For more information on FAST.X refer to [FAST.X](#) on page 266.

Use the `symdev list` command with the `-external` option to list information about eDisks (external spindles).

FAST.X is limited to the following:

- Virtual provisioning encapsulation of Data Domain eDisks for ProtectPoint.
 - The `symdev list` command includes a new `Encapsulated Device Flags` field to indicate when an encapsulated device is a Data Domain device.
 - A device status of *Not Ready* for a Data Domain device indicates that the device is *Not Ready* because it is specified as the target of an image refresh.
- External provisioning for XtremIO, VNX, VMAX and some third party arrays.

List external spindles

Syntax

To list external spindles, use the following syntax:

```
symdev [-sid SymmID] [-offline] [-v] [-resv | -pgr]
      [-wnn | -wnn_encapsulated [-detail] | -wnn_non_native] [-all]
] [-all]
  list [-FA <#|ALL> [-P <#>] |
      -SA <#|ALL> [-scsi] [-fibre] [-P <#>]]
  . . .
  [-technology <EFD | FC | SATA>]
  [-internal | -external | -encapsulated [-limited]]
  list pd [-FA <#|ALL> [-P <#>] |
      -SA <#|ALL> [-scsi] [-fibre] [-P <#>]]
  . . .
  [-technology <EFD | FC | SATA>]
  [-internal | -external | -encapsulated [-limited]]
symdev [-sid SymmID] [-offline]

  list [-DA <#|ALL> | -DX <#|ALL>] -space [-cyl]
  [-spindle]

  list [-DA <#|ALL>] [-interface <#|ALL>]
  [-disk <#|ALL>] [-hyper <#|ALL>] [-firstport]
  [-spindle] [-internal]
  list [-DX <#|ALL>] [-hyper <#|ALL>]
  [-spindle] [-external]
```

Options

- DX**
Lists devices on a specific external director (DX) or all DX directors. The existing DA option only shows devices on an internal DA director.
- internal**
Lists devices on internal spindles. TDEV, VDEV and DLDEV devices are not considered either internal or external and are not returned by this option, and is abbreviated as `-int`.
- external**
Lists devices on external spindles. TDEV, VDEV and DLDEV devices are not considered either internal or external and are not returned by this option, and is abbreviated as `-ext`.
- encapsulated**
Lists encapsulated devices, and abbreviated as `-enc`.
- limited**
Lists devices that are geometry limited. These two options can also be used to report which TDEVs are considered encapsulated or geometry limited.

NOTE: The `symdisk list` command also provides options for listing external spindle information. For more information, also refer to [External spindle information](#) on page 365.

List RAID group information

Description

The `symdev` command indicates whether a rebuild or copy is in progress for a member of a device's RAID groups. The `RAID Group Information` section includes an additional indicator under the `Status` column, which displays when a specific member is undergoing one of the following:

- Rebuild — (R)
- Member copy — (C)
- Failed state — (F)

Examples

To list if a rebuild or copy is in progress for a member of the RAID groups for device 00050, enter:

```
symdev show 00050 -sid 123
```

Sample output

```
. . .
Device Service State      : Normal
. . .
RAID Group Information
{
. . .
RAID Group Service State  : Normal
. . .
  {
    Device : 00050
    {
      -----
      Spindle  Disk      Hyper      Member      Mir  Attr
              DA :IT    Num  Cap (MB)  Num Status
      -----
      30E      07A:D5    1    61         1 RW (C)    1  N/A
      3FE      09A:D5    1    61         2 RW        1  N/A
    }
  }
}
```

Report PowerPath device status

Description

Use the `-ppi` option with the `symdev list` command to display PowerPath mount status of array devices.

Syntax

To display PowerPath mount status of array devices, use the following syntax:

```
symdev list -ppi [-sid SymmID]
```

Examples

To list PowerPath device status on array 789, enter:

```
symdev list -ppi -sid 789
```

Sample output

```
Symmextrix ID: 000198700789
      P O W E R P A T H   D E V I C E   S T A T U S
Device Last Used           Mounted  Hostname      Process name
-----
0001F  06/29/17 08:46:51      Yes   Host_A        RedoLog
00341  06/29/17 08:47:52      No    Host_B        DBWrite
```

Report devices with non-native WWN

Description

Use the `-wwn_non_native` option with the `symdev list` command and `symdev show` command to list or show devices with Device External Identity set to non-native WWN.

Syntax

To list devices with Device External Identity set to non-native WWN, use the following syntax:

```
symdev [-sid SymmID] list -wwn_non_native
```

To show a device with Device External Identity set to non-native WWN, use the following syntax

Syntax - symdev show

To show a device with Device External Identity set to non-native WWN, use the following syntax:

```
symdev show -wwn_non_native <WWN>
```

Examples

To list all the devices on array 085 with a non-native WWN, enter:

```
symdev -sid 085 list -wwn_non_native
```

To show an external WWN for device 60000970000196801476533030314245, enter:

```
symdev show -wwn_non_native 60000970000196801476533030314245
```

Sample output

For `symdev list`:

```
Symmetrix ID: 000197100085
      Device Name           Device
-----
Sym  Physical              Config      Attr Non-Native WWN
-----
0003D Not Visible          RDF2+TDEV  60000970000196801476533030313831
0003E Not Visible          RDF2+TDEV  60000970000196801476533030313832
0005E Not Visible          RDF2+TDEV  60000970000196801476533030313833
```

For `symdev show`:

```
Symmetrix ID: 000194901137
      Device Physical Name      : Not Visible
      Device Symmetrix Name     : 01AD4
```

```

Device Serial ID      : N/A
Symmetrix ID         : 000194901137

. . .

Product Revision     : 5977
Device WWN           : 60000970000194901137533031414434
Device Emulation Type : FBA

. . .

Device External Identity
{
  Device WWN         : 60000970000196801476533030314245
}

. . .

```

Change the device state

Syntax

To change the device state, set or reset the hold bit on a device, or relabel a device, use the following syntax:

```

symdev -sid SymmID [-devs <<SymDevStart>:<SymDevEnd> | <SymDevName>
                    [,<<SymDevStart>:<SymDevEnd> | <SymDevName>>...]]
[-noprompt] [-rp] [-celerra] [-star]
  rw_enable      [-SA <#|ALL>[-P <#>]]

  write_disable [-SA <#|ALL>[-P <#>]]

  ready

  not_ready

  relabel

  hold

  unhold        [-symforce]

```

Change device state using filename

Syntax

To change the device state of multiple devices listed in a file, use the following syntax:

```

symdev -sid SymmID -file FileName
[-noprompt] [-celerra] [-star]
  rw_enable      [-SA <#|ALL> [-P <#>]]

  write_disable [-SA <#|ALL> [-P <#>]]

  ready

  not_ready

  relabel

  hold

  unhold        [-symforce]

```

View device service states

Description

The `symdev list` and `symdev list pd` commands provide options that return data that is limited to a specific service state.

Options

`-service_state`

Returns list of devices one of the following states:

- Normal
- Failed
- Degraded — Indicates if one or more mirror positions of the protected device are not ready.

To list service states not in a particular state use "not" such as, `notfailed`.

Examples

To list all devices on array 097 that are in a `normal` state, enter:

```
symdev list -service_state normal -sid 097
```

to list all devices on array 097 that are `notfailed` service state, enter:

```
symdev list pd -service_state notfailed -sid 097
```

Device emulation

All host I/O transactions with an array of disk devices are managed by the array operating system, which runs in the array I/O subsystem (channel directors and disk directors). Because each of the physical disks are indirectly seen as part of the I/O protocol, array devices are presented to the host with the following configuration or emulation attributes:

- Each device has N cylinders. The number is configurable (blocks ÷ 960).
- Each cylinder has 15 tracks (heads).
- Each device track in a fixed block architecture (FBA) has 64 blocks of 64K. (For non-FBA operating systems, the blocks are recognized without regard to the number of bytes.)

List device encryption flags

Examples

To see the type of device encryption for array 056, enter:

```
symdev show -sid 056
```

Sample input

```
. . . .
Device WWN                : 600009700BC7233831110020000000E9
Device ID Type            : Mobility
Device Emulation Type     : FBA
Device Encryption        : None
. . . .
Device WWN                : 600009700BC7233831110020000000EA
Device ID Type            : Mobility
Device Emulation Type     : FBA
Device Encryption        : EffEncryptCapable
. . . .
Device WWN                : 600009700BC7233831110020000000EB
```



```

Device ID Type      : Mobility
Device Emulation Type : FBA
Device Encryption   : EffEncrypted
. . .
Device WWN          : 600009700BC7233831110020000000EC
Device ID Type      : Mobility
Device Emulation Type : FBA
Device Encryption   : RekeyInProg
. . .
Device WWN          : 600009700BC7233831110020000000ED
Device ID Type      : Mobility
Device Emulation Type : FBA
Device Encryption   : NoValidKey
. . .

```

List device emulation types

Examples

To return a list of configured devices by emulation type for array 3139, enter:

```
symdev list -inventory -sid 3139
```

Sample input

```

Symmetrix ID: 000192603139
  Device Config      FBA      CKD3390  CKD3380  AS400  CELERRA
-----
Unprotected          256      N/A      N/A      N/A      N/A
2-Way Mir            4040     N/A      N/A      N/A      N/A
RAID-5               483      N/A      N/A      N/A      N/A
RAID-6               515      8        N/A      N/A      N/A
TDEV                 496      N/A      N/A      N/A      N/A
DLDEV                2381     N/A      N/A      N/A      N/A
RDF1+Mir              8        N/A      N/A      N/A      N/A
RDF1+R-6             16       N/A      N/A      N/A      N/A
RDF1+DLDEV           1275     N/A      N/A      N/A      N/A
RDF2+Mir              40       N/A      N/A      N/A      N/A
RDF2+R-5              8        N/A      N/A      N/A      N/A
RDF2+R-6              8        N/A      N/A      N/A      N/A
RDF21+R-5            16       N/A      N/A      N/A      N/A
RDF21+TDEV           16       N/A      N/A      N/A      N/A
BCV                   10       N/A      N/A      N/A      N/A
2-Way BCV Mir         5        N/A      N/A      N/A      N/A
BCV+R-5               5        N/A      N/A      N/A      N/A
BCV+R-6               5        N/A      N/A      N/A      N/A
2-Way DRV Mir         10       N/A      N/A      N/A      N/A
VDEV                  213     N/A      N/A      N/A      N/A
BCV+TDEV              32       N/A      N/A      N/A      N/A

```

Show device details

Examples

To details for device 8423 on array 584, enter:

```
symdev show 8423 -sid 584
```

Sample output

```
. . .
Mirror Configuration Information
{
  Mirror Number           : 1
  Mirror Type             : RAID-6
  Mirror Status           : Ready           (RW)
}
RAID Group Information
{
  Mirror Number           : 1
  RAID Type               : RAID-6
  Device Position         : Primary
  Protection Level        : 6+2
  Disk Group Number       : 1
  Disk Group Name         : DISK_GROUP_001
  Engine Number           : Multiple
  RAID Group Service State : Normal
  Number of Failing Members : 0
  Member Information:
  {
    Device : 8423
    {
      -----
      Spindle   Disk       Hyper      Member   Mir   Attr
              DA  :IT    Num Cap(MB) Num Status
      -----
      12FC      09C:D1   10      699     1 RW    1   N/A
      120C      07C:D1    9      699     2 RW    1   N/A
      111C      05C:D1    9      699     3 RW    1   N/A
      198C      07D:D0    9      699     4 RW    1   N/A
      1A7C      09D:D0    9      699     5 RW    1   N/A
      960       05B:C1   10      699     6 RW    1   N/A
      438       10A:C1   10      699     7 RW    1   N/A
      1914      06D:D1   10      699     8 RW    1   N/A
    }
  }
}
```

Show clone state flags

The following CLI list and show commands report the Clone State Flags:

- `symdev show`
- `symdev list -v`
- `sympd list -v`
- `sympd show`
- `symdg show ld`
- `symdg list ld -v`
- `symcg show ld`

Show disk geometry details

Examples

To show disk geometry for device 0016 on array 516, enter:

```
symdev show 0516 -geometry -sid 516
```

Sample output

NOTE: If the array-wide setting is turned on, the geometry at the individual device level can be defined which overrides the array-wide setting. In addition, for non-FBA devices, the output displays N/A with the `-v` and the `-geometry` options.

```
Device Physical Name      : /dev/sdd
Device Symmetrix Name    : 00516
Device Serial ID        : 1600016000
Symmetrix ID            : 000190300516
Attached BCV Device     : N/A
Attached VDEV TGT Device : N/A
Vendor ID               : EMC
Product ID              : SYMMETRIX
Product Revision        : 5977
Device WWN              : 60060480000190300516533030303136
Device Emulation Type   : FBA
Device Defined Label Type : N/A
Device Defined Label    : N/A
Device Sub System Id    : 0x0001
Device Block Size       : 512
Device Capacity
{
  Cylinders      :      4400
  Tracks         :      66000
  512-byte Blocks :    4224000
  MegaBytes      :       2063
  KiloBytes      :    2112000
}
Effective Device Geometry:  User Defined
{
  Sectors/Track   :       128
  Tracks/Cylinder :       15
  Cylinders       :       2000
  512-byte Blocks :    3840000
  MegaBytes       :       1875
  KiloBytes       :    1920000
}
Device Configuration   : RDF1                (Non-Exclusive Access)
Device is WORM Enabled : No
Device is WORM Protected : No
SCSI-3 Persistent Reserve: Disabled
Dynamic Spare Invoked  : No
Dynamic RDF Capability : None
STAR Mode              : No
STAR Recovery Capability : Sync_Tgt
STAR Recovery State    : Inactive
Device Service State   : Normal
Device Status          : Not Ready          (NR)
Device SA Status       : Ready             (RW)
Front Director Paths (2):
```

Output display

The Effective Device Geometry field has the following possible values:

- User Defined — Indicates the user has defined geometry for this device.
- Native — Indicates the current device geometry is the same as the native geometry for the array.
- Array wide emulation — Indicates that the array-wide flag for FBA geometry emulation is set to enabled and the effective geometry shown in the output is derived from this setting.

List DATA and SAVE devices

Syntax

To list DATA and SAVE pool device details, use the following syntax:

```
symcfg [-sid SymmID] [-offline] [-mb | -gb]
list [-savedevs [-fba] [-ckd3390] [-ckd3380] [-as400]
      [-nonpooled]
      [-devs [SymDevStart:SymDevEnd | SymDevName]
           [, SymDevStart:SymDevEnd | SymDevName...]]
list [-datadev [-fba] [-nonpooled]
      [-devs [SymDevStart:SymDevEnd | SymDevName]
           [, SymDevStart:SymDevEnd | SymDevName...]]
```

NOTE: When the `symcfg list` command is used to return pool information, the returned data includes pools that are currently associated with the group, or pools that have been disassociated from the group but still have the group's data.

Options

- all**
Includes both enabled and disabled devices in the calculations of `Free Tracks` and `Full %`, and the `Full %` is based on the `Usable Tracks`. Otherwise, only enabled devices are included. Without the `-all` option, the `Free Tracks` field includes free tracks from all enabled tracks, and the `Full %` field is based on the enabled tracks.
- devs**
Filters for specific devices or device range.
- nonpooled**
Filters for nonpooled devices.

Examples

To return a list of DATA devices for array 237, enter:

```
symcfg list -sid 237 -datadev
```

List thin device information

Description

In addition to listing thin devices on an array, use the `symcfg list -dev` command to monitor the progress of space reclamation. The `Reclaiming` status indicates the thin device is in the process of being reclaimed and the `Total Allocated Tracks (%)` adjusts as space is reclaimed.

Syntax

To list thin devices in an array, use the following syntax:

```
symcfg [-sid SymmID] [-offline] [-mb | -gb]
list [-tdev
      [-devs SymDevStart:SymDevEnd | SymDevName
           [, SymDevStart:SymDevEnd | SymDevName...]]
      [-pool PoolName | [-fba] [-ckd3390]
           [-bound | -unbound]]
      [-sg SgName] [-detail | -tier]]
```

Options

- tdev
Lists the thin devices in the system.
- bound
Lists thin devices bound to a pool.
- unbound
Lists thin devices that are unbound.

Examples

To list the thin devices on array 341, enter:

```
symcfg list -tdev -sid 341
```

Sample output

NOTE: The columns **Total Allocated Tracks (%)** displays the percentage of the entire thin device allocated.

```
Symmetrix ID: 000194900341
Enabled Capacity (Tracks) :      785088
Bound Capacity (Tracks) :      816420
      S Y M M E T R I X   T H I N   D E V I C E S
-----
Sym      Bound      Flgs      Total      Total      Compressed
Pool Name EMPT      Tracks      Allocated  Size/Ratio
                                     Tracks (%)  Tracks (%)
-----
0035B example2  F..B      33000      32796 99    32976 0
0035C test1     F..B      33000      1092 3    1092 0
00A68 testing_341 F..B      33000      13428 41    13428 0
00B1D xyz       F..B      14400      36 0    36 0
00B1E example2  F..B      48000      12 0    12 0
. . .
Total
Tracks      89557755      9560808 63 9560808 0
Legend:
Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
(M)ultipool : X = multi-pool allocations, . = single pool allocation
(P)ersistent Allocs : A = All, S = Some, . = None
S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
          D = Deallocating, R = Reclaiming, C = Compressing,
          N = Uncompressing, . = Unbound
```

NOTE: Show details about thin pools on page 351 provides information on compressing and uncompressing thin devices.

List thin device tier allocations

Examples

To list tier allocations for thin devices on array 341, enter:

```
symcfg list -tdev -sid 341 -tier
```

Sample output

Lists both technology and disk locations for the tiers.

```
Symmetrix ID: 000194900341
      S Y M M E T R I X   T H I N   D E V I C E S
-----
Total      Tier
```

Sym	Emul	Total Tracks	Allocated Tracks (%)	Flag LT	Protection	Name
01078	FBA	33000	0 0	IF	RAID-5(3+1)	FC_R5_VPTier
01079	FBA	33000	0 0	--	-	-
0107A	FBA	66000	144 0	IS	RAID-6(6+2)	SATA_R6_VPTier
010EC	FBA	33000	132 0	IF	RAID-6(14+2)	FC_R6_VPTier
010EE	FBA	33000	0 0	--	-	-
010F9	FBA	1500	20 0	X-	Unprotected	EXT_VPTIER
01128	FBA	8010	24 0	IF	RAID-6(14+2)	Finance_VPTier
011A2	FBA	33000	216 1	--	-	[OutOfTier]
Total		240510	536 0			

Legend:
Disk (L)ocation:
I = Internal, X = External, M = Mixed, - = N/A
(T)echnology:
S = SATA, F = Fibre Channel, E = Enterprise Flash Drive, M = Mixed, - = N/A

List thin devices in a pool

Examples

To display a list of thin devices in pool HR_THIN_R5, enter:

```
symcfg -sid 343 list -tdev -pool HR_THIN_R5
```

Show thin BCV devices with dynamic SRDF capability

Examples

To show device 15F9 on array 397, enter:

```
symdev show 15F9 -sid 397
```

Sample output

Device Configuration displays device 15F9 as BCV+TDEV and Dynamic RDF Capability as RDF1_OR_RDF2_Capable.

```
Device Physical Name      : Not Visible
Device Symmetrix Name    : 015F9
Device Serial ID        : N/A
Symmetrix ID            : 000194900397
Number of RAID Groups   : 0
Attached VDEV TGT Device : N/A
Vendor ID               : EMC
Product ID              : SYMMETRIX
Product Revision        : 5977
Device WWN              : 60000970000194900397533031354639
Device Emulation Type   : FBA
Device Defined Label Type : N/A
Device Defined Label    : N/A
Device Sub System Id    : 0x0016
Cache Partition Name    : DEFAULT_PARTITION
Bound Pool Name         : N/A
. . .
Device Configuration    : BCV+TDEV
Device is WORM Enabled  : No
Device is WORM Protected : No
SCSI-3 Persistent Reserve : Enabled
Dynamic Spare Invoked   : No
Dynamic RDF Capability   : RDF1_OR_RDF2_Capable
```

NOTE:

Bound Pool Name will be reported as N/A for an encapsulated device.

List data allocations across multiple pools

Options

`-detail`

Lists thin devices and all the pools where tracks are allocated to these devices.

Examples

```
symcfg -sid 397 -devs 1620:1630 list -tdev -detail
```

Sample output

```
Symmetrix ID: 000195700397
Enabled Capacity (Tracks) : 2884812
Bound Capacity (Tracks) : 20700
S Y M M E T R I X T H I N D E V I C E S
-----
```

Sym	Bound Pool Name	Flgs ESPT	Total Tracks	Pool Subs (%)	Pool Allocated Tracks (%)	Compressed Size/Ratio Tracks (%)
1623	testCust	F..B	14700	13	2400 16	1692 30
	test_pool	----	-	-	252 2	0 0
1625	-	8...	1500	0	0 0	0 0
1626	-	8...	1500	0	0 0	0 0
1627	-	8...	1500	0	0 0	0 0
162E	testCust	F..B	6000	10	4800 80	1200 75
Total			25200	1	7452 30	2892 61

```
Legend:
Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
(S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
(P)ersistent Allocs : A = All, S = Some, . = None
S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
          D = Deallocating, R = Reclaiming, C = Compressing,
          N = Uncompressing, . = Unbound
```

Show details about thin pools

Options

`-detail`

Shows details about a Thin Provisioning pool.

`-all`

Shows the disabled devices in the pool.

Examples

To show details about thin pool DG1_FBA_PL on array 086, enter:

```
symcfg show -pool DG1_FBA_PL -thin -detail -sid 086
```

Sample output

The Pool State displays as Enabled, Disabled, or Balancing. If the Thin Provisioning pool is enabled and the write balancing feature is turned on, the Pool State displays as Balancing. A thin device status of Reclaiming indicates the thin devices are in the process space reclamation.

Any thin devices bound to a specified pool are listed along with thin devices that have allocated tracks within the pool but are not bound to that pool, displays in Other Thin Devices with Allocations in this Pool.

```
. . .
```

```

Pool Name           : DG1_FBA_PL
Pool Type           : Thin
. . .

Enabled Devices(208):
{
-----
Sym      Usable      Alloc      Free Full FLG Device
Dev      Tracks      Tracks      Tracks (%) S  State
-----
1FF30    274680      96         274680  3  .  Enabled
1FF31    274680      96         274680  3  .  Enabled
. . .
Tracks   1056000      40332      1054464  3
}

```

No Thin Devices Bound to Device Pool DG1_FBA_PL

Other Thin Devices with Allocations in this Pool (3):

```

{
-----
Sym      Bound      Total      Pool      Compressed
Pool Name Tracks      Tracks (%) Size/Ratio
-----
00003 -          33000      32796  99      0  0
00004 -          49500      7536  15      0  0
. . .
Tracks   82500      40332  48      0  0
}

```

Legend:

Enabled devices FLG:

(S)hared Tracks : X = Shared Tracks , . = No Shared Tracks

Bound Devices FLG:

S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,

D = Deallocating, R = Reclaiming, C = Compressing,

N = Uncompressing, F = FreeingAll, . = Unbound,

NOTE: The **FLG S** flag field indicates whether there are shared allocations on each device as noted in the above legend.

Verify thin and DATA device states

Description

More than one type of device state may be queried. All devices must be in any of the requested states for the verification to succeed. Normally, the command executes a single poll and exits with status 0 (all devices in requested state) or nonzero status indicating which devices were not ready. A message, such as, "One or more of the specified devices are not in the requested state", returns.

Syntax

To verify the state of thin or DATA devices, use the following syntax:

```

symcfg -sid SymmID [-i Interval] [-c Count]
[-pool PoolName |
-g DgName |
-sg SgName |
-cg CgName |
-devs SymDevStart:SymDevEnd | SymDevName[, SymDevStart:SymDevEnd | SymDevName...]
symcfg verify -tdev
[-bound | -binding | -allocating | -deallocating |
-unbound | -unbinding | -reclaiming]
symcfg verify -datadev
[-draining | -drainwait | -disabled | -enabled |
-deactivated | -nonpooled | -balancing]

```


Options

-devs

Specifies the status of single or multiple devices.

-pool

Returns all thin or DATA devices within a pool and must be in one of the requested states to pass verification.

-i

Specifies the polling rate. The default minimum interval value used for calculating the action is 15 seconds. A message displays at each poll. If the **-i** option is used without the **-c** option, the command loops infinitely, exiting only when all devices are in the correct state (or when manually interrupted).

-c

Specifies the numbers of polls to be executed. The command terminates at the end of the count or when all devices are in the correct state.

Verify device usage

Examples

To check once to see if all of the DATA devices in a pool (`devpool`) on array `1234` are disabled, enter:

```
symcfg -sid 1234 verify -datadev -pool devpool -disabled
```

To continually check the status of a thin device (`5AC`) on array `1234` and exit the loop when it enters a bound state, enter:

```
symcfg -sid 1234 verify -tdev -dev 5ac -bound -I 5
```

Migrated devices

Use the `symdev list` and `symdev show` commands to provide information about devices that are changed or virtualized during a migration.

List changed devices after migration

Options

-identity_set

Displays the devices whose effective identity changed after a migrate operation.

Examples

To display the devices with identity change on array `207`, enter:

```
symdev -sid 207 -identity_set list
```

Sample output

```
Symmetrix ID: 000192600207
-----
Device Name          Directors          Device
-----
Sym Physical          SA :P DA :IT Config          Attribute      Sts      Cap
(MB)
-----
0477 /dev/sddv        ***:* 06E:C1 2-Way Mir      N/Grp'd (M) RW      2063
```

List virtualized devices

Options

-identity

Displays the effective identity information of devices that are virtualized during a migrate operation.

Examples

To list identity information of migrated device 0477 on array 207, enter:

```
symdev -sid 207 -identity show 0477
```

Sample output

```
Device Physical Name      : Not Visible
Device Symmetrix Name     : 00477
Device Serial ID         : N/A
Symmetrix ID             : 000192600207
Number of RAID Groups    : 1
Attached BCV Device       : N/A
Attached VDEV TGT Device : N/A
Vendor ID                : EMC
Product ID               : SYMMETRIX
Product Revision         : 5825
Device WWN                : 60000970000192600207533030343737
Device Emulation Type    : FBA
Device Defined Label Type: N/A
Device Defined Label      : N/A
Device Sub System Id     : 0x0001
Cache Partition Name     : DEFAULT_PARTITION
Device Block Size        : 512
Device Capacity
  Cylinders              :          960
  Tracks                 :         14400
  512-byte Blocks       :        1843200
  MegaBytes              :           900
  KiloBytes              :         921600
Effective Device Information
  Device WWN              : 60000970000192600306533030313732
  Front Director Paths (12): (See note below)
-----
  DIRECTOR  PORT              LUN
-----
  Type Num  Sts  VBUS TID SYMM Host
-----
  FA N/A  07E:3  RW   000 00  247  N/A
  FA N/A  07F:3  RW   000 00  247  N/A
Geometry: User Defined
{
  Sectors/Track      :          64
  Tracks/Cylinder    :          15
  Cylinders           :          200
  512-byte Blocks    :        192000
  MegaBytes           :           94
  KiloBytes           :         96000
Device Configuration : Unprotected
Device is WORM Enabled : No
Device is WORM Protected : No
SCSI-3 Persistent Reserve: Disabled
Dynamic Spare Invoked : No
Dynamic RDF Capability : RDF1_OR_RDF2_Capable
STAR Mode             : No
STAR Recovery Capability : None
STAR Recovery State   : NA
Device Service State  : Normal
Device Status         : Ready (RW)
Device SA Status      : N/A (N/A)
Host Access Mode      : Active
Device Tag(s)         : None
```

```
Front Director Paths (21):
```

PdevName	POWERPATH	DIRECTOR	PORT		LUN			
	Type	Type Num	Sts	VBUS	TID	SYMM	Host	
Not Visible	N/A	FA	07E:1	RW	000	00	247	N/A
Not Visible	N/A	FA	07F:1	RW	000	00	247	N/A

```
Remote Copy Device Information
```

```
Remote Copy Session Name      : N/A
Session State                  : Created
Pull                           : True
Offline Copy                   : False
Differential Copy              : Disabled
Background Copy                : Disabled
Incremental Copy               : False
Donor Update                   : True
Federated Live Migration       : True
Starting Block                 : 0
Total Blocks                   : 192000
Percent Copied                 : N/A
Remote Devices [1]
```

Array:Devv	WWN	Starting Block
000187990007:0012E	60000970000192600306533030313732	0

Filter list for device data

Description

The `symdev list` and `symdev list pd` commands return a list of devices configured in one or more arrays connected to the controlling host. The `pd` qualifier returns only physical devices, and must appear after the list action on the command line, and is valid with all the `symdev list` options.

Options

-all

Lists all configured devices.

Syntax

To list device information with filter options, use the following syntax:

NOTE: For more information on device list filter options, refer to the *EMC Solutions Enabler SYMCLI Command Reference*.

```
symdev [-sid SymmID] [-offline] [-v] [-resv | -pgr]
        [-wwn | -wwn_encapsulated [-detail]] [-all]

list [ -FA <#|ALL> [-p <#>] |
      -SA <#|ALL> [-scsi] [-fibre] [-p <#>]]
      [-devs SymDevStart:SymDevEnd | SymDevName
        [, SymDevStart:SymDevEnd | SymDevName...]]
      [-cap <#> [-captype <mb | cyl>]] [-N <#>]
      [-vcm | -aclx] [-held] [-gige]
      [-ficon] [-escon] [-dldev] [-dd] [-dm]
      [-noport|-firstport|-multiport] [-bcv|-nobcv|-drv]
      [-meta] [-nomember]
      [-spare] [-dynamic] [-worm] [-vdev]
      [-savedev [-nonpooled]] [-raids]
      [-disk_group DskGrpNum | name:DskGrpName]
      [-rg] [-sec_raid] [-unprotected] [-sec_unprotected]
      [-raid1] [-sec_raid1] [-bcv_emulation]
      [-raid5 [-protection <3+1 | 7+1>]]
      [-raid6 [-protection <6+2 | 14+2>]]
```

```

[-sec_raid5 [-sec_protection <3+1 | 7+1>]]
[-sec_raid6 [-sec_protection <6+2 | 14+2>]]
[-emulation fba|ckd|ckd3390|ckd3380|as400|celerra]
[-star_mode] [-star_sync_target] [-star_async_target]
[-reserved | -noreserved] [-cyl] [-geometry_set]
[-service_state [not]degraded | [not]failed |
                 [not]normal]
[-tdev [-bound | -unbound]] [-datadev [-nonpooled]]
[-migr_tgt] [-pinned] [-host_passive] [-identity_set]
[-identity [-detail]] [-sg SgName]
[-rp [Production | Replica | Internal]]
[-technology <EFD | FC | SATA>]
[-host_cache] [-dif1] [-as400_gk] [-fast]
[-internal | -external | -encapsulated [-limited]]
[ -notrdf | [-rdfg RdfGrpNum [-rdfa]
[-R1] [-R2] [-R21] [-half_pair] [-dup_pair]]]
  [-insg | -notinsg
[-host_id <compatible | native>]

```

Filter device list by director

- SA #**
Lists devices by a specific front-end director number.
- p #**
Lists devices by a specific front-end director port number.
- scsi**
Lists devices mapped to SCSI front-end directors.
- fibre**
Lists devices mapped to Fibre front-end directors.

Filter device list by director port mapping

- multiport**
Lists devices mapped to more than one front-end director port.
- noport**
Lists devices not mapped to any front-end director port.
- firstport**
Lists first port information for devices with more than one port mapping schemas.

If none of these options are specified, then devices with all director-port relationships return.

Filter device list by metadvice

- meta**
Lists meta member devices.
- nonmember**
Lists only meta head devices.

Filter device list by SRDF devices

- R1**
Lists all SRDF R1 devices.
- R2**
Lists all SRDF R2 devices.
- R21**
Lists all SRDF R21 devices. The -R21 option is not currently supported in conjunction with the `-dynamic` option.
- rdfg**
Lists only devices belonging to the specified SRDF group number (`RdfGrpNum`).

- rdfa**
Lists all SRDF/Asynchronous capable devices.
- half_pair**
List all devices that are not paired with an SRDF device. Existing half pair devices can result from an SRDF/Star failover scenario, a half_deletepair operation, or a configuration change.
- dup_pair**
Lists all devices that are paired with the same SRDF type. Existing duplicate pair devices can result from a SRDF/Star failover scenario or a configuration change.
- notrdf**
Lists all dynamic capable devices that are not SRDF devices. Use this option to identify non-SRDF devices that have dynamic SRDF capability. In addition, use this option with the -R1, -R2, and -R21 device options to include non-RDF devices in the list along with these specified devices.

NOTE: Refer to the *EMC Solutions Enabler SRDF Family CLI Product Guide* for detailed information on SRDF device types and failover configurations.

Filter device list by RAID information

- raids**
Lists RAID-S devices by by RAID group number. The RAID option lists the same information. When viewing device output for RAID-5 devices, the total device capacity (`Device Capacity`) is displaying the total logical capacity, and the hyper capacity (`Hyper Num Cap`) is displaying, for each member, the physical capacity which includes all device space—available and unavailable.

Filter device list by technology type

- technology**
Lists the drive technology (FC, EFD, and SATA) of the primary local back-end storage of the device.

Filtering device list by DA, interface, disk, or hypervolume

- da, -dx, -interface, -disk, -disk_group, -hyper**
Lists specified DAs, DXs, interfaces, disks, disk groups or hypervolumes. These options default to ALL if not specified.

Filtering by storage group

- sg**
Lists devices by specified storage group. This option defaults to ALL if not specified.

Filtering by devices in data migration

- dm**
Lists array devices that are in data migration session.

Filter device list by DIF1 attribute

- dif1**
Lists device DIF1 attribute status. The Data Integrity Field (DIF) is a setting on a device that is relevant to an Oracle environment and all hosts that support the DIF protocol. If the DIF1 attribute is set on a device, it displays as TRUE or FALSE.

Restrictions

- DIF1 attribute can only be set on FBA devices.
- DIF1 attribute can be set on both standard provisioned (thick) and virtually provisioned (thin) host accessible devices, setting the DIF1 attribute request on any internal devices will be rejected.
- Devices must be unmapped while resetting the DIF1 attribute.
- Devices can be mapped to only fiber front end directors (no iSCSI or FCoE) while setting the DIF1 attribute.
- If a device has DIF1 attribute already set, it can be mapped only to fiber front end directors (no iSCSI or FCoE).

- If creating a meta device, both the head and members need to have the DIF1 attribute in the same state (either all set or all reset).
- DIF1 attribute needs to be set before requesting reset. If the reset request is for a device range and if any one of the device has the DIF1 attribute not set, then the request fails.
- DIF1 attribute can not be set on DATA and SAVE devices.
- Devices can only have the RDB_Checksum attribute or the DIF1 attribute set , but not both.
- The DIF1 flag can not be set on a device having active DCS.
- Devices can only have the ACLX attribute or the DIF1 attribute set, but not both.
- There is no relationship between the DIF1 attribute and replication. Both source and target devices of any replication can have their own DIF1 setting.

Device list capacity output options

-cyl

Lists device capacity in cylinders rather than the default of megabytes (MB).

-pd

Lists only the host visible devices (pdevs).

Device external locks

Solutions Enabler uses *device external locks* in the array to lock pairs during replication operations (such as Open Replicator, TimeFinder, and SRDF operations).

NOTE: Use the release lock action only if it is determined that the device lock was forgotten and there are NO other operations in progress to the specified devices (local or remote). Locks are typically of short duration (one second to an hour or so). However, it is critical to be able to recognize when a device lock held by certain applications (such as an SRDF action) are longer duration locks.

List devices with external lock

Examples

To list a range of devices (0000:000A) that have a device external lock, enter:

```
symdev list -sid 870 -devs 0000:000A -lock
```

Release devices with external lock

Examples

To release the locks held by the configuration server for device 204 on array 343, enter:

```
symdev release -sid 343 devs 204
```

```
Within Symmetrix unit 000190300343, release lock for device 204 (y/[n]) ? y
Device lock is held by a migration session <654>. The session should be aborted to release
the lock.
```

NOTE: Aborting the session, if successful, releases all locks held by the session including device locks.

Disk-level data

The configuration database file maintains low-level configuration and status information for each disk on every array that are accessible from the host, including external spindles (eDisks). Use the `symdisk list` and `symdisk show` commands, to list all the available disks, or specify individual disks to obtain configuration and status information.

List and show disk information

Syntax

To display disk details, use the following syntax:

```
symdisk [-sid <SymmID>] [-offline] [-cyl | -mb | -gb | -tb]
        [-disk_group <DskGrpNum | name:<DskGrpName> | ALL>]
        [-all] [-failed]
list [-spindle [-internal]] [-isspare] [-nospare]
     [-v [-hypers] [-spare_info] [-gaps]]
     [-DA <# | ALL>] [-interface <# | ALL>]
     [-tid <# | ALL>]

list [-spindle
     [-external [-detail] [-encapsulated [-free]]]]
     [-v [-hypers] [-gaps]]
     [-DX <# | ALL>]
symdisk [-sid <SymmID>] -external -spindle -paths
        [-spid <SpindleID> | -DX <# | ALL> [-port <# | ALL>]]
list -detail
list [-offline]
symdisk [-sid <SymmID>] [-offline] [-cyl | -mb | -gb | -tb]
list -dskgrp_summary [-v]
     [-disk_group <DskGrpNum | name:<DskGrpName> | ALL>]
     | -internal | -external]
list -dskgrp_summary -by_engine [-v | -detail]
     [-disk_group <DskGrpNum | name:<DskGrpName> | ALL>]
     | -internal | -external]
show <DiskAddress> [-gaps_only]
symdisk -sid <SymmID> [-offline] [-cyl | -mb | -gb | -tb]
show -spid <SpindleID> [-gaps_only]
        show -wwn <ExternalWWN> [-gaps_only]
```

Options

-engine

Lists attributes, counts and capacities relevant to Internal Disk Group/Engine Spindle Groups.

-cyl | -mb | -gb | -tb

Indicates how to report disk capacity.

List SAS drives

Description

Serial Attached SCSI (SAS) disk drives in 300 GB, 450 GB and 600 GB capacity versions are supported. This includes the Cobra-D SAS 300, 450, and 600 GB disks drives (2.5" drive using 3.5" carrier and SAS/FC paddle card).

NOTE: For Disk Groups that contain internal spindles with different form factors (e.g. 2.5" and 3.5"), the Form Factor label displays as Mixed.

Options

-v

Lists the SAS technology type in the Technology field. The Form Factor field displays 2.5, 3.5, or N/A to indicate the Enterprise Flash Drive.

List disk gaps

Options

- gaps**
Used with the `symdisk list` command, and lists any gaps found on the disk.
- spare_info**
Used with the `symdisk list` command, and shows which disk the spare disk is substituting for, if it is invoked.
- gaps_only**
Used with the `symdisk show` command, and lists any gaps found on the disk without listing all of the hyper information.

Reported actual disk capacity vs. rated disk capacity

The `Actual Disk Capacity` lists the physical disk capacities in MBs and GBs, where the MB is defined as (1024 x 1024) bytes and the GB is defined as 1024 MBs (or 1024 x 1024 x 1024 bytes).

The `Rated Disk Capacity` is based on the MB being defined as (1000 x 1000 bytes), the GB being defined as 1000 MBs (or 1000 x 1000 x 1000 bytes), and the TB defined as 1000 GBs.

List disk groups

Options

Used with the `symdisk list` command.

- by_diskgroup**
Lists the disks by disk group number.
- disk_group**
Limits disks that belong to a specified disk group. Disk groups can be specified by group number (`DskGrpNum`) or group name (name: `DskGrpName`). The ALL option returns disk information for all disk groups.
- dskgrp_summary**
This option displays summary information for disk groups. Without the `-v` option the disk group summary information is displayed in a table format. With the `-v` option the disk group information in expanded format (one field per line) is displayed, information regarding spare disks is not included.

NOTE: The `-dskgrp_summary` option is not valid with the `-hypers`, `-spare_info`, `-gaps`, `-isspare` and `-failed` options.

List disk group summary

Examples

To list disk group summary for array 584, enter:

```
symdisk list -dskgrp_summary -sid 584
```

NOTE: The command line accepts `-dskgrp_summary` abbreviated as `dskg`.

Sample output

```
Symmetrix ID: 000197100584
  Disk Group                Disk                Hyper                Capacity
-----
 Num  Name                  Flgs  Speed  Size  Size  Total  Free
      Cnt  LT  (RPM) (MB)  (MB)  (MB)  (MB)  (MB)
```


1	DISK_GROUP_001	67	IS	7200	1823565	3561	122178855	0
2	DISK_GROUP_002	67	IF	15000	1823565	3561	122178855	0
512	DISK_GROUP_512	1	X-	N/A	N/A	Any	204801	201114
513	DISK_GROUP_513	1	X-	N/A	N/A	Any	204801	102400
Total							244767312	303514

Output with `syndisk list verbose (-v)` option:

```

Disk Group                : 1
Disk Group Name           : DISK_GROUP_001
Disks Selected            : 67
Disk Location              : Internal
Technology                 : SATA
Speed (RPM)               : 7200
Form Factor                : 3.5
Disk Size (MB)            : 1823565
Rated Disk Size (GB)      : 2000
Total Group Capacity (MB) : 122178855
Free Group Capacity (MB)  : 23005664
Max Hypers Per Disk       : 512
Hyper Size (MB)           : 3561
.....

```

List disk groups by engine

Options

`-by_engine`

Lists attributes, counts and capacities specific to internal disk group/engine spindle groups.

`-detail`

Expands the list to include the disk group name.

Examples

To list disk groups by engine for array 132, enter:

```
syndisk list -dskgrp_summary -by_engine -sid 132
```

NOTE: The command line accepts `-dskgrp_summary` abbreviated as `dskg`, and `-by_engine` abbreviated as `-by_eng`.

Sample output

```

Symmetrix ID: 000197800132

-----
Disk              Hyper    Usable Capacity    Spare Coverage
-----
Grp  Eng  Cnt  Flgs Speed Size  Disk Total Total Avail
-----
      LT  (RPM) (MB)  (%)  (MB)  Disk (%) Disk (%)
-----
  1   1   18  IE    0  28470  16  89  29153150  2  12  2  100
  1   2   18  IE    0  28470  16  89  29153150  2  12  2  100
  2   1   25  IE    0  14235  24  96  21864765  1   4  1  100
  2   2   25  IE    0  14235  24  96  21864765  1   4  1  100
-----
Total                                80  92  102035830

Legend:
Disk (L)ocation:
  I = Internal, X = External, - = N/A
(T)echnology:
  S = SATA, F = Fibre Channel, E = Enterprise Flash Drive, - = N/A

```

Output field description

- Usable Capacity — Reports in units of MB/GB/TB, based on `-mb/-gb/-tb` flag setting. The usable capacity is also reported in units of disks and percentage (relative to the total capacity for the internal disk group/engine spindle group) for internal disk groups.

- **Total Spare Coverage** — Reports in units of Disk and percent (% for internal disk groups/engine spindle groups. The Disk-based units correspond to the Disk Size and Max Hypers Per Disk reported in the `-diskgrp_summary -v` output. The Total Spare Coverage values include all spare capacity, including unavailable capacity (i.e. failed disks), within each internal disk group/engine spindle group and the percentage is relative to the Usable capacity.
- **Available Spare Coverage** — Reports in units of Disk and percent (% for internal disk groups/engine spindle groups. The Disk-based units correspond to the Disk Size and Max Hypers Per Disk reported in the `-diskgrp_summary -v` display. The Available Spare Coverage values are the amount of spare capacity currently available for sparing and this percent is relative to the Total Spare Coverage value.

Output with `syndisk list (-detail)` option:

```
Symmetrix ID: 000197100584
```

Disk Group		Disk	Hyper	Usable Capacity			Spare Coverage						
Num	Name	Eng	Cnt	LT	Flgs	Speed	Size	Disk (%)	Total (MB)	Total Disk (%)	Total Disk (%)	Avail (%)	
1	DISK_GROUP_001	1	50	IS		7200	113972	48 96	87530496	2 4	2 100		
1	DISK_GROUP_001	2	50	IS		7200	113972	48 96	87530496	2 4	2 100		
1	DISK_GROUP_001	3	100	IS		7200	113972	96 96	175060992	4 4	3 75		
2	DISK_GROUP_002	4	50	IF		15000	113972	49 98	107589568	1 2	0 0		
2	DISK_GROUP_002	5	25	IF		15000	113972	24 96	43765248	1 4	1 100		
2	DISK_GROUP_002	6	100	IF		15000	113972	98 98	178708096	2 2	2 100		
512	DISK_GROUP_512	7	1	X-		N/A	Any	N/A N/A	204800	N/A N/A	N/A N/A		
Total										363	97	680389696	

Legend:
 Disk (L)ocation:
 I = Internal, X = External, - = N/A
 (T)echnology:
 S = SATA, F = Fibre Channel, E = Enterprise Flash Drive, - = N/A

List spare physical disks

Options

Use the following options with `-DA`, `-interface`, or `-tid` options to list spare disk information on specific directors, disk interfaces, or SCSI target IDs.

-isspare

Lists information about spare physical disks (spindles), if a spare disk is invoked against a failed disk.

-v -spare_info

Lists information about the failed disk that has been replaced. The `-v` option must be specified with this option. This display option is not applicable for permanent sparing.

-gaps

Lists gaps found between hypers on the disk, as the hypers are listed. To see a short list of only the gap information, do not specify the `-hypers` option on the `syndisk list` command, or specify the `-gaps_only` option on the `syndisk show` command.

NOTE: The gap sizes provided by this report are approximate values. The report can be used as a general guide to the location and size of free space gaps, but it may not be accurate down to the last cylinder.

-nospare

Lists only internal disks (spindles) that are capable of being covered by a spare, but currently are not. This option is not valid for external spindles (external spares do not exist) and no external spindles display when this option is specified.

Examples

To list information spare disk information about failed disk, enter:

```
syndisk list -v -spare_info
```

Sample output

Lists a spare disk that is invoked against disk 16B:C0

```
...
Spare Disk           : False
Director            : DF-16B
Interface           : C
Target ID           : A
Disk Group Number   : 0
Vendor ID           : SEAGATE
Product ID          : SX3146807FC
Product Revision    : CH146LF
Serial ID           : 3HY9CQVK
Disk Blocks         : 0
Block Size          : 512
Actual Disk Blocks  : 286749475
Total Disk Capacity (MB) : 0
Free Disk Capacity (MB) : 0
Actual Disk Capacity (MB) : 140014
Hypers              : 0
Spare Disk          : True
Failed Director     : DF-15A
Failed Interface    : C
Failed Target ID    : 0
```

List spindle information

Options

-spindle

Lists each disk by Spindle ID and includes the director number for each spindle.

-spid *Spindle_ID*

Lists physical disks by Spindle ID.

Examples

To list the spindle IDs for array 306, enter:

```
symdisk list -spindle -sid 306
```

Sample output

```
Symmetrix ID           : 000194900306
Disks Selected         : 64

Spindle  Dir  Vendor      Type      Hypr  Total  Capacity (MB)
-----  --  -
          Dir  Vendor      Type      Hypr  Total  Free      Actual
-----  --  -
32E      08A  SEAGATE    T146155   14    139028  120506    139028
16B35    08A  SEAGATE    T146155   19    139028  118758    139028
1D83     08B  SEAGATE    T146155   15    139028  123071    139028
22C34    08B  SEAGATE    T146155   15    139028  124751    139028
. . .
-----  --  -
Totals                    7785556  6975177  7785556
```

List spindle information for physical disk

Syntax

-spid *Spindle_ID*

Specifies spindle ID.

Examples

To list spindle information for spindle 1B935 on array 306, enter:

```
symdisk show -spid 1B935 -sid 306
```

NOTE: The `sympd list` command with the `-spindle` option also provides spindle ID information. [List array device information on page 337](#) provides more information.

Sample output

NOTE: This is the same output format as `symdisk list -v` (verbose).

```
Symmetrix ID           : 000194900306
Director              : DF-8A
Interface             : C
Target ID             : 1
Spindle ID            : 1B935
External WWN          : 60000970000195700233533030333132
External Array ID     : 0001957000233
External Device Name   : 031B
Disk Group Number     : 001
Disk Group Name       : DISK_GROUP_001
Technology            : FC
Speed (RPM)           : 15000
Form Factor           : N/A
Vendor ID             : EMC
Product ID            : Symmetrix
. . .
```

NOTE: The `symdisk list` and `show` commands for spindles identifies remote Data Domain appliances when reporting on virtualized disks and LUNs on external arrays. If applicable, `DataDomain` displays in the `Product ID` field.

List spindle information for disk group

Options

`-spindle`

Lists each disk by Spindle ID and includes the director number for each spindle.

Examples

To list spindle IDs for disk group 004 for array 306, enter:

```
symdisk list -disk_group 004 -spindle -sid 306
```

Sample output

```
Symmetrix ID           : 000194900306
Disks Selected         : 8
Disk Group             : 4
Disk Group Name       : DISK_GROUP_004
Technology            : FC
Speed (RPM)           : 15000
```

Spindle	Dir	Vendor	Type	Hypr	Total	Capacity (MB)	
						Free	Actual
32A	07A	SEAGATE	HUC4515	52	418710	380208	418710
B7	08A	SEAGATE	HUC4515	52	418710	380208	418710
5C9	07B	SEAGATE	HUC4515	53	418710	379729	418710
1D01	08B	SEAGATE	HUC4515	53	418710	364605	418710
. . .							
Total					3349684	3009503	3349684

External spindle information

FAST.X attaches external disks (eDisks) to a VMAX arrays, and makes the eDisk's capacity available to the VMAX array as an external spindle.

Use the `symdisk show` command to list the spindle IDs for eDisks.

 **NOTE:** For more information on eDisks, refer to [FAST.X](#) on page 266

View external spindle information

Syntax

To list external spindle information, use the following syntax:

```
symdisk show -wwn wwn
```

Output fields for external spindle information

- `Disk Location` — Shows `Internal` for internal disks, and `External` for external disks.
- `Disk Service State` — Shows the availability of the eDisk. The `Failed Disk State` is reported as part of the `Disk Service State`, showing states of *Normal* or *Failed* for internal disks. For external disks the `Disk Service State` is shown as *Failed* if there are no network paths available to the eDisk (neither active nor passive), *Degraded* if there are paths from only one of its supporting DX directors to the eDisk (either active or passive), and *Normal* if there is at least one active and one passive network path available from both supporting DX directors to the eDisk.
- `Technology` — Shows as `N/A` for external disks.
- Shows if an eDisk is encapsulated or not.
- Shows the number of external paths, available to the external LUN, to reach the eDisk.

Events and Logs

This chapter describes how to configure, manage, and query VMAX events and logs.

Topics:

- [Events and logs overview](#)
- [Log option configuration](#)
- [Array event monitoring using SYMCLI](#)
- [Common audit log](#)
- [Daemon log files](#)

Events and logs overview

SYMCLI and SYMAPI normally log significant events and actions to a daily log file. On UNIX, the log file has the following pathname:

```
/var/symapi/log/symapi-yyyymmdd.log
```

On Windows, the log file has the following pathname:

```
C:\Program Files\EMC\Symapi\log\symapi-yyyymmdd.log
```

File name date format:

- *yyyy* — year
- *mm* — month
- *dd* — day

The log displays the following items about each event:

- Time tag of the event occurrence
- Process ID (PID)
- Source of the event (application name)
- Related (internal) API function call
- Name of the specific operation or event
- Variable event field that describes the event or error in detail

NOTE:

By default, SYMAPI log files are retained forever. You can change this to automatically remove the files after a set amount of time by modifying the SYMAPI_LOGFILE_RETENTION option in the `options` file. It is recommended that you configure a log file retention time to conserve disk space.

The format for the retention option is SYMAPI_LOGFILE_RETENTION = *NN*. Where *NN* is the number of days to retain the SYMAPI log file.

Default value: 0 (retain forever)

Valid values: 5 (5 days) - 1825 (5 years)

Log option configuration

The following table lists the options to customize how and where events are logged in your VMAX environment.

Table 23. Log file configuration options

Logging option	Configuration description
Enable/disable logging	<p>You can disable logging by setting the environment variable <code>SYMCLI_NOLOGGING</code> to 1. For example, to disable logging on UNIX (C shell), enter:</p> <pre>setenv SYMCLI_NOLOGGING 1</pre> <p>To turn logging back on, enter:</p> <pre>unsetenv SYMCLI_NOLOGGING</pre>
Log file name	<p>You can change the name and path of the default log file.</p> <p>For example, to change the log file name on UNIX (C shell), use the following form:</p> <pre>setenv SYMCLI_LOG filename</pre> <p>To turn daily log files back on, enter:</p> <pre>unsetenv SYMCLI_LOG</pre>
Allow undated SYMAPI log file names	<p>You can allow the creation of undated SYMAPI log files by setting the environment variable <code>SYMAPI_DATED_LOGFILE_NAME</code> in the options file to <code>disable</code>.</p> <p>To allow undated log file names, set:</p> <pre>SYMAPI_DATED_LOGFILE_NAME=DISABLE</pre> <p>To reenab date, set:</p> <pre>SYMAPI_DATED_LOGFILE_NAME=ENABLE</pre>
Change date formats	<p>You can change date formats in the log entries by setting the environment variable <code>SYMAPI_LOGFILE_DATE_FORMAT</code> in the options file to <code>FORMAT2</code>. This formats the date as <code>yyyy-mm-dd</code>.</p> <p>To change the date format, set:</p> <pre>SYMAPI_LOGFILE_DATE_FORMAT=FORMAT2</pre> <p>To change the date format back to the default (<code>mm/dd/yyyy</code>), set:</p> <pre>SYMAPI_LOGFILE_DATE_FORMAT=FORMAT1</pre>

Table 23. Log file configuration options (continued)

Logging option	Configuration description
Log file configuration options	<p>You can change the format of optional fields within each log record by setting the environment variable SYMAPI_LOGFILE_FORMAT in the options file. Zero or more of the following optional fields can be included:</p> <ul style="list-style-type: none"> • <code>pid</code> — include the process ID. • <code>tid</code> — include the thread ID. • <code>userid</code> — include the user ID (useful with User Authorization is enabled). • <code>activityid</code> — include the activity ID. <p>The default value is <code>pid tid</code>. To see the log file format, create an entry in the options file with values separated by spaces. For example:</p> <pre>SYMAPI_LOGFILE_FORMAT = userid pid tid activityid</pre>
Log file retention	<p>Log file retention is an option set in the options file to specify the number of days to retain the log files:</p> <ul style="list-style-type: none"> • Maximum value: 1825 (or 5 years) • Minimum value: 6 • Alternative value: 0 (setting it to zero maintains the log file forever. This is the default for everything except the service processor whose default is 30.) <p>For example, to change the log file Retention to 60 days, enter:</p> <pre>SYMAPI_LOGFILE_RETENTI ON=60</pre> <p>NOTE:</p> <p>When the log file retention option is set, it overrides all existing log file Retention values including those set for the service processor during initialization. The default for the service</p>

Table 23. Log file configuration options (continued)

Logging option	Configuration description
	<p>processor is 30 days. Setting this options file value will override that default. In addition, if you have accumulated many years of log files prior to setting this option, when the option is set, it deletes those log files that do not meet the specified criteria.</p>

Array event monitoring using SYMCLI

The `symevent` command allows an administrator to monitor or track events within a VMAX array that may affect its operation. In some cases, reported events represent conditions that have already been repaired. This command allows monitoring of the array for all reported events.

Use **`symevent`** to:

- `monitor` — This action runs the command in the foreground, polling the array for new events every interval in seconds, either until the iteration count is satisfied or the program is stopped. Monitoring can be restricted to report events of certain severity (warnings, errors, or fatal events).
- `list` — This action examines the history of events stored on the array, for those events that meet the requested criteria, such as events recorded during a certain time period, or based upon the reporting director.

Monitoring events

Examples

To poll for and report on all events, on all locally-connected arrays, every 15 seconds, continuously, enter:

```
symevent monitor
```

To poll for and display events with a severity of *warning* or greater every 15 seconds for a 50-second period, enter:

```
symevent monitor -sid 0207 -i 15 -c 50 -warn
```

Sample output

For array 0207:

```
Symmetrix ID: 000192600207
Detection time      Dir      Src  Category      Severity      Error Num
-----
Mon Dec 29 21:26:04 2015 DF-7A  Symm  Communication Warning      0x001a
    The Symmetrix Service Processor could not complete a Call Home for service
```

List events

Options

`-start / -end`

Lists events occurring during a certain time period.

-DIR

Limits reporting to specific director activity.

Examples

To retrieve a verbose list of the events which have occurred on the specified array between 12 p.m. and 3 p.m. today, enter:

```
symevent list -sid 0207 -start 12:00 -end 15:00 -v
```

Common audit log

Data is written to a common audit file during VMAX control operations. The common audit log correlates activity from all hosts into one file that is stored in the Symmetrix File System (SFS).

The `symaudit` command enables the filtering of the common audit log file for a specified VMAX array. The audit log resides on the array with a capacity of 40 MB. Once the 40 MB limit is reached, the log starts to overwrite itself. There is no maintenance for this file, unless records need to be captured before the circular 40 MB space recycles.

NOTE: For details on filtering parameter values refer to the `symaudit` command in the Dell EMC Solutions Enabler CLI Reference Guide.

Use the following actions to parse or monitor the audit log contents:

show

Displays details about the audit log itself for a specific array, including the total range of records, the date/time range, and the starting record number.

list

Lists details about the requested records in either a brief or verbose format. The range of records to extract can be filtered by one or all of the following:

- record number
- record count
- date/time
- functional area
- control action
- vendor ID
- application ID
- hostname
- username
- device acted upon

monitor

Monitors the array for new audit log data in realtime.

NOTE: If Embedded Management is configured on the VMAX array, when running the `symaudit list` and `symaudit show` commands the host name listed in the output is the nodename identified by the NAT gateway, and not the internal identity of the eManagement Guest. For more information on eManagement refer to the EMC VMAX3 Family Product Guide for VMAX 100K, VMAX 200K, VMAX 400K with HYPERMAX OS.

Show audit log for single array

Examples

```
symaudit show -sid 0207
```

Sample output

```
A U D I T   L O G   D A T A
Symmetrix ID       : 000192600207
Starting date      : 12/17/2008 21:42:57
Ending date       : 01/02/2009 15:30:35
Starting record number : 1
Ending record number   : 1737
Total record count   : 1737
```

List audit log for specified time period

Examples

```
symaudit list -sid 0207 -start_date 01/02/2009:12:00:00 -end_date 01/02/2009:12:15:00 -v
```

Sample output

```
A U D I T   L O G   D A T A
Symmetrix ID       : 000192600207
Record Number      : 1663
Records in Seq     : 2
Offset in Seq      : 1
Time              : 01/02/16 12:11:50
Vendor ID         : EMC Corp
Application ID     : SYMCONFIGURE
Application Version : 8.2.0.302
API Library       : SDK
API Version       : T8.2.302.2 (Edit Level: 907)
Host Name        : api1051.lss.
OS Name          : LINUX
OS Revision       : 2.6.9-11.E
Client Host      :
Process ID       : 00027021
Task ID          : 00000002
Function Class    : CfgChg
Action Code       : Release
Text             : STARTING a Device Reservation 'RELEASE'. Owner=m; ReserveID=1;
Comment="t";
Username          : H:api1051\root
Activity ID       : SE69f82805ea
Record Number     : 1664
Records in Seq   : 2
Offset in Seq    : 2
Time            : 01/02/16 12:11:50
Vendor ID       : EMC Corp
Application ID  : SYMCONFIGURE
Application Version : 8.2.0.302
API Library     : SDK
API Version     : T8.2.302.2 (Edit Level: 907)
Host Name      : api1051.lss.
OS Name        : LINUX
OS Revision    : 2.6.9-11.E
Client Host    :
Process ID     : 00027021
Task ID       : 00000002
Function Class : CfgChg
Action Code    : Release
Text          : Devices: [ 0020 ]
Username      : H:api1051\root
Activity ID   : SE69f82805ea
Record Number : 1665
Records in Seq : 1
Offset in Seq  : 1
Time          : 01/02/16 12:11:51
Vendor ID     : EMC Corp
```

```

Application ID      : SYMCONFIGURE
Application Version : 8.2.0.302
API Library        : SDK
API Version        : T8.2.302.2 (Edit Level: 907)
Host Name          : api1051.lss.
OS Name            : LINUX
OS Revision        : 2.6.9-11.E
Client Host        :
Process ID         : 00027021
Task ID            : 00000002
Function Class     : CfgChg
Action Code        : Release
Text               : Device Reservation 'RELEASE' SUCCEEDED. ReserveID=1;
Username           : H:api1051\root
Activity ID        : SE69f82805ea
<. . .>

```

Output field descriptions:

- **Record Number** — The current record number.
- **Records in Seq** — Total number of records requested.
- **Offset in Seq** — Offset number from the first record requested.
- **Time** — Date and time the record was entered.
- **Vendor ID** — ID of the vendor whose application logged the record.
- **Application ID** — ID of the application that logged the record.
- **Application Version** — Application version number.
- **API Library** — Name of the SYMAPI library the application ran against.
- **API Version** — Version of the SYMAPI.
- **Host Name** — Name of the host that logged the record.
- **OS Name** — Operating system on which the host is running.
- **OS Revision** — Operating system revision number.
- **Client Host** — Any SYMCLI client communicating with the SYMAPI server.
- **Process ID** — ID of the process that logged the record.
- **Task ID** — ID of the task.
- **Function Class** — Class name of the SYMAPI functional area.
- **Action Code** — Name of the SYMAPI control action associated with an audit log entry.
- **Text** — Details of the given entry.
- **Username** — Identifies the user that generated the log entry.
- **Activity ID** — A randomly generated ID that uniquely identifies this action.

Custom audit log activity ID

By default, Solutions Enabler generates a random activity ID to identify each session. This ID appears in the audit log entries for that session. Since the default ID is random, filtering audit log entries based on the ID is difficult.

You can specify a custom activity ID, making it easier to filter audit log entries.

An optional argument `-actid <Activity ID>`, allows you to set a custom activity ID on all operations performed as part of that CLI command.

When you use the `-actid <Activity ID>` argument, entries in the audit log for that command are tagged with the specified activity ID prefixed with "U_".

User-defined activity IDs must meet the following requirements:

- Maximum of 14 characters (not including the automatic prefix).
- Include only alphanumeric characters.

Underscore (`_`), and hyphen (`-`) characters are allowed.

If you try to define an Activity ID with more than 14 characters, the operation fails, and an error is displayed.

All active SYMCLI commands, with the exception of `symaudit` command, accept the `-actid` argument. For example:

```
symqos -sid 237 -cp -name cptest -devs 410:411,1170 addall -actid
CPTest1
```

Output of symaudit queries in verbose mode report the user-defined activity ID. For example:

```
symaudit list -sid 237 -v -activity_id U_CPTest1
```

```

      A U D I T L O G D A T A
Symmetrix ID      : 000190300237
  Record Number   : 711815
    Records in Seq : 2
.
.
.
  Username       : H:dldv0181\root
  Activity ID    : U_CPTest1
  Record Number   : 711816
    Records in Seq : 2
.
.
.

```

Daemon log files

Each Solutions Enabler daemon has two log files to record daemon errors and other significant conditions.

Daemon log files are stored in:

```
<SYMAPI_HOME>/log/storXXXX.log0
```

```
<SYMAPI_HOME>/log/storXXXX.log1
```

where *storXXXX* is the name of the daemon. The daemons are:

- storsrvd (SYMAPI Server Daemon)
- storevntd (Event Daemon)
- storgnsd (GNS Daemon)
- storrdfd (RDF Daemon)
- storapid (Base Daemon)
- storstpd (STP Daemon)
- storwatchd (Watchdog Daemon, UNIX only)

Logging alternates between the two files, switching to the other file each time the maximum size specified by the daemon's LOGFILE_SIZE parameter is reached. Each daemon writes to the .log0 file until its size exceeds that specified in the LOGFILE_SIZE option, at which point it switches to the .log1 file. It switches back to .log0 under the same conditions.

The daemon_options file includes options that allow you to customize logging for a particular daemon, including the type of log file (wrap or dated), the maximum size for wrap files, the amount of time to retain dated log files, and log file permissions.

Defaults

- LOGFILE_SIZE parameter is 1 MB.
- Daemon log files use wrap mode. To change the log files to dated mode, use the LOGFILE_TYPE option in the daemon_options file. In dated mode, the daemon log file format is *daemon_name-yyyyymmdd.log*. A new dated log file is started every day on the first write after 12:00 a.m.
- Dated log files are retained for three days. You can change this to automatically remove the files after a set amount of time by modifying the LOGFILE_RETENTION option in the daemon_options file.

The daemon_options file includes options that allow you to further customize logging for a particular daemon.

The following table shows the general logging configuration options you can use to customize the Solutions Enabler daemon log files. For details on the syntax and values, refer to the <SYMAPI_HOME>/config/daemon_options file.

Table 24. Options for managing daemon log files

Option	Description
LOGFILE_TYPE	Specifies the style of logging for the daemon. Possible values:

Table 24. Options for managing daemon log files (continued)

Option	Description
	<p>WRAP (default) - Two log files are maintained: storxxx.log0 and storxxx.log1. Logging switches to the other file each time the maximum size specified by the LOGFILE_SIZE parameter is reached.</p> <p>DATED - A separate log file is used for each day: storxxx-YYYYMMDD.log. There are no size limits on these files. Dated log files are retained for the number of days specified by the LOGFILE_RETENTION parameter.</p>
LOGFILE_SIZE	<p>Used for wrapping log files.</p> <p>Specifies the maximum number of KBs to write before switching to the other file of the pair.</p> <p>Default: 1000 (KB)</p>
LOGFILE_RETENTION	<p>Used for dated log files.</p> <p>How many days to retain old log files.</p> <p>Default: 3 (days)</p> <p>Valid values: A number of days greater than 0.</p>
LOGFILE_PERM	<p>Specifies the permissions on any newly created log files. Possible values:</p> <p>rw (default) - Anyone can read or write (UNIX mode of rwxrwxrwx).</p> <p>r - The owner (root) can read/write.</p>

Event daemon

Solutions Enabler also provides an asynchronous event daemon that can be used to monitor and manage VMAX events. The event daemon (`storevntd`) provides the services required to monitor the status of VMAX storage environments from third-party enterprise management frameworks. The following targets are supported:

- SNMP
- File on disk
- System logger on the host
- Unix syslog service
- Windows event log
- The syslog listener across the network (bypasses the syslog service (calls) on the local host and directly sends events/traps to this remote listener).

For more information on enabling and configuring the event daemon, refer to the *EMC Solutions Enabler Installation Guide*.

XML Structured Output

This chapter describes the XML output option of the SYMCLI.

Topics:

- [XML structured output overview](#)
- [XSLT: XML data transformations](#)
- [Element-based XML](#)
- [Set XML mode with SYMCLI](#)

XML structured output overview

The XML (Extensible Markup Language) output option provides a mechanism to facilitate the automated processing of SYMCLI output data. XML is a deterministic parsing tool that eases parsing of output data, providing a functional advantage over screen-scraping tools like awk or Perl. The XML industry standard is based on the experience of SGML and is endorsed by the World Wide Web Consortium. Detailed information on XML may be found at: <http://w3.org/XML/>.

XML has the look and feel of HTML, as it employs the same tag-based syntax. However, XML uses tags to delimit data—as opposed to defining the data as with HTML—allowing the document author to specify the tags most applicable to the given application. For the SYMCLI, tags represent the physical and logical structures within the VMAX array and its environments.

When XML mode is utilized, the data returned is identical to that of the standard output, but "marked-up" with tags. These tags enable individual pieces of data to be readily called upon by name. In addition, they provide a definitive way to express the relationship between different objects, an advantage over the standard CLI display output.

XSLT: XML data transformations

Many tools are available to query, filter, retrieve, and format specific information stored in complex XML files. Among these, eXtensible Stylesheet Language Transforms (XSLT) is a particularly useful and widely available technology. While using XML will result in less ambiguous, more robust scripts, XSLT will make the information presented in XML accessible to the more familiar plain text-based scripting techniques. To introduce you to XSLT, a directory containing several examples of the types of queries that can be performed on XML data using XSLT is provided. The examples are designed only to provide a brief introduction to the power and usefulness of XSLT, and can help ease the transition to XML.

Element-based XML

Solutions Enabler provides element-based XML that describes data in a hierarchical manner by using the notion of parent and children. An element can have several different content types. It can have element content (child element), a mixed content containing both text and child element, a simple content containing text only, or an empty content carrying no information. An element can also have attributes. These additional content types would allow users to modify the data structures in a fairly flexible manner. On the other hand, an attribute is used to provide additional information about an element. An attribute is, in general, used to store the metadata describing the data that stored in XML. Although data can be stored in attributes, it is best practice to store data in child elements.

Set XML mode with SYMCLI

To use XML mode with SYMCLI, an environment variable or a command line option can be used. To use the environment variable to globally set the command output to XML or standard use the following syntax:

```
SYMCLI_OUTPUT_MODE = xml_element|standard
```



When the environment variable output mode is set to `xml` or `xml_element`, commands that do not support XML output generate a runtime error message. To override this behavior set the command line `-output` option to a value of `standard`. This allows execution of a given command in standard mode.

XML output using SYMCLI

Syntax

To override the current environment variable setting and set the output style for a single command, use the following syntax:

```
<SymcliCommand> -output <xml|xml_element|standard>
```

NOTE:

The `-output` flag is not found in `-help` or man pages because of its wide scope and usage.

Options

`xml_element`

Returns the output of all commands in element-based XML tags.

`standard`

Returns the output of all commands in the default output format.

Examples

```
symcfg list -out xml or symcfg list -out xml_element
```

Sample output

Note that a new element tag `<Symm_Info>` is added to the XML data shown below to store general VMAX data.

```
<?xml version="1.0" standalone="yes" ?>
<SymCLI_ML>
  <Symmetrix>
    <Symm_Info>
      <symid>000190102055</symid>
      <attachment>Local</attachment>
      <model>VMAX3-100K</model>
      <microcode_version>5977</microcode_version>
      <cache_megabytes>32768</cache_megabytes>
      <devices>683</devices>
      <physical_devices>94</physical_devices>
    </Symm_Info>
  </Symmetrix>
  <Symmetrix>
    <Symm_Info>
      <symid>000190300215</symid>
      <attachment>Local</attachment>
      <model>VMAX3-100K</model>
      <microcode_version>5977</microcode_version>
      <cache_megabytes>32768</cache_megabytes>
      <devices>239</devices>
      <physical_devices>2</physical_devices>
    </Symm_Info>
  </Symmetrix>
  <Symmetrix>
    <Symm_Info>
      <symid>000190300237</symid>
      <attachment>Remote</attachment>
      <model>VMAX3-100K</model>
      <microcode_version>5977</microcode_version>
      <cache_megabytes>16384</cache_megabytes>
      <devices>3756</devices>
```



```
<physical_devices>0</physical_devices>
</Symm_Info>
</Symmetrix>
<Symmetrix>
  <Symm_Info>
    <symid>000190300343</symid>
    <attachment>Remote</attachment>
    <model>VMAX3-100K</model>
    <microcode_version>5977</microcode_version>
    <cache_megabytes>32768</cache_megabytes>
    <devices>751</devices>
    <physical_devices>0</physical_devices>
  </Symm_Info>
</Symmetrix>
</SymCLI_ML>
```

To maintain consistent element names in all SYMCLI commands, some tag names are redefined. For example, in current SYMCLI command, different names exist to describe Symmetrix identification number, such as id, symmetrix, or symid. A new consistent tag name is defined across all SYMCLI commands in the element-based XML output.

Migrating Data

This section describes:

- how to use Non-Disruptive Migration (*syndm*) to migrate data
 - from VMAX arrays running Enginuity 5876 to VMAX3 Family and VMAX All Flash arrays running HYPERMAX OS 5977 Q3 2016 SR and higher
 - from VMAX3 Family and VMAX All Flash arrays running HYPERMAX OS 5977 Q3 2016 SR and higher to PowerMax arrays running PowerMaxOS
- how to use Open Replicator (*symrcopy*) to migrate data between EMC arrays or third-party arrays.

Chapters include:

Topics:

- [Non-Disruptive Migration](#)
- [Open Replicator](#)

Non-Disruptive Migration

This chapter describes the Non-Disruptive Migration (NDM) feature and the SYMCLI command `symdm` command used to migrate data:

- from Engenuity 5876 to VMAX3 Family and VMAX All Flash arrays running HYPERMAX OS 5977 Q4 2016 SR and higher.
- from Engenuity 5876 to PowerMax arrays running PowerMaxOS 5978.
- from HYPERMAX OS 5977 Q4 2016 SR or higher to PowerMax arrays running PowerMaxOS 5978.

For details on supported releases, please consult the *Dell EMC SRDF Interfamily Connectivity Guide*.

Topics:

- [Non-Disruptive Migration overview](#)
- [Non-Disruptive Migration for PowerMaxOS 5978 overview](#)
- [Non-Disruptive Migration operational restrictions and state reference](#)
- [Non-Disruptive Migration control operations](#)
- [List Non-Disruptive Migration session status](#)
- [Non-Disruptive Migration examples](#)

Non-Disruptive Migration overview

Non-Disruptive Migration (NDM) provides a method for migrating data from a source array to a target array without application host downtime across a metro distance, typically within a data center. For NDM array operating system version support, please consult the [NDM support matrix](#), or the *SRDF Interfamily Connectivity Guide*.

If regulatory or business requirements for DR (disaster recovery) dictate the use of SRDF/S during migration, contact Dell EMC for required ePacks for SRDF/S configuration.

The NDM operations involved in a typical migration are:

- Environment setup – Configures source and target array infrastructure for the migration process.
- Create – Duplicates the application storage environment from source array to target array.
- Cutover – Switches the application data access from the source array to the target array and duplicates the application data on the source array to the target array.
- Commit – Removes application resources from the source array and releases the resources used for migration. Application permanently runs on the target array.
- Environment remove – Removes the migration infrastructure created by the environmental setup.

Some key features of NDM are:

- Simple process for migration:
 1. Select storage group to migrate.
 2. Create the migration session.
 3. Discover paths to the host.
 4. Cutover or readytgt storage group to VMAX3 or VMAX All Flash array.
 5. Monitor for synchronization to complete.
 6. Commit the migration.
- Allows for inline compression on VMAX All Flash array during migration.
- Maintains snapshot and disaster recovery relationships on source array, but are not migrated.
- Allows for non-disruptive revert to source array.
- Allows up to 50 concurrent migration sessions.
- Requires no license since it is part of HYPERMAX OS.
- Requires no additional hardware in the data path.

The following graphic shows the connections required between the host (single or cluster) and the source and target array, and the SRDF connection between the two arrays.

Figure 12. Non-Disruptive Migration zoning



The App host connection to both arrays uses FC, and the SRDF connection between arrays uses FC or GigE .

The migration controls should be run from a control host and not from the application host. The control host should have visibility to both the source array and target array.

The following devices and components are not supported with NDM:

- CKD devices, IBM i devices
- eNAS data
- ProtectPoint, FAST.X relationships and associated data

Environmental requirements for Non-Disruptive Migration

The following configurations are required for a successful data migration:

Array configuration

- The target array must be running HYPERMAX OS 5977.811.784 or higher. This includes VMAX3 Family arrays and VMAX All Flash arrays.
- The source array must be a VMAX array running Enginuity 5876 with required ePack (contact Dell EMC for required ePack).
- SRDF is used for data migration, so zoning of SRDF ports between the source and target arrays is required. Note that an SRDF license is not required, as there is no charge for NDM.
- The NDM RDF group is configured with a minimum of two paths on different directors for redundancy and fault tolerance. If more paths are found up to eight paths will be configured.
- If SRDF is not normally used in the migration environment, it may be necessary to install and configure RDF directors and ports on both the source and target arrays and physically configure SAN connectivity.

Host configuration

- The migration controls should be run from a control host and not from the application host.
- Both the source and the target array have be visible to the controlling host that runs the migration commands.

Pre-migration rules and restrictions for Non-Disruptive Migration

In addition to general configuration requirements of the migration environment, the following conditions are evaluated by Solutions Enabler prior to starting a migration.

- A Storage Group is the data container that is migrated, and the following requirements apply to a storage group and its devices:

- Storage groups must have masking views. All devices within the storage group on the source array must be visible only through a masking view. The device must only be mapped to a port that is part of the masking view.
- Multiple masking views on the storage group using the same initiator group are only allowed if port groups on the target array already exist for each masking view, and the ports in the port groups are selected.
- Storage groups must be parent or standalone storage groups. A child storage group with a masking view on the child storage group is not supported.
- Gatekeeper devices in the storage group are not migrated to the target array.
- Devices must not be masked or mapped as FCoE ports, iSCSI ports, or non-ACLX enabled ports.
- The devices in the storage group cannot be protected using RecoverPoint.
- Devices may not be in multiple storage groups.
- For objects that may already exist on the target array, the following restrictions apply:
 - The names of the storage groups (parent and/or children) to be migrated must not exist on the target array.
 - The names of masking views to be migrated must not exist on the target array.
 - The names of the initiator groups (IG) to be migrated may exist on the target array. However, the aggregate set of host initiators in the initiator groups used by the masking views must be the same. Additionally the effective ports flags on the host initiators must be the same for both source and target arrays
 - The names of the port groups to be migrated may exist on the target array, provided that the groups on the target array appear in the logging history table for at least one port.
- The status of the target array must be as follows:
 - If a target-side Storage Resource Pool (SRP) is specified for the migration that SRP must exist on the target array.
 - The SRP to be used for target-side storage must have enough free capacity to support the migration.
 - The target side must be able to support the additional devices required to receive the source-side data.
 - All initiators provisioned to an application on the source array must also be logged into ports on the target array.
- Only FBA devices are supported (Celerra and D910 are not supported) and the following restrictions apply:
 - Cannot have user geometry set, mobility ID, or the BCV attribute.
 - Cannot be encapsulated, a Data Domain device, or a striped meta device with different size members.
 - Must be dynamic SRDF R1 and SRDF R2 (DRX) capable and be R1 or non-RDF devices, but cannot be R2 or concurrent RDF devices, or part of a Star Consistency Group.
- Devices in the storage group to be migrated can have TimeFinder sessions and/or they can be R1 devices. The migration controls evaluates the state of these devices to determine if the control operation can proceed.
- The devices in the storage group cannot be part of another migration session.

Unsupported devices

The following list provides an overview of unsupported devices, and device restrictions:

- Only IBM i devices are supported (CKD is not supported) with the following restrictions:
 - Cannot have user geometry set, mobility ID, or the BCV attribute.
 - Cannot be encapsulated, a Data Domain device, or a striped meta device with different size members.
 - Must be dynamic SRDF R1 and SRDF R2 (DRX) capable and be R1 or non-RDF devices, but cannot be R2 or concurrent RDF devices, or part of a Star Consistency Group.
- Devices in the storage group to be migrated can have TimeFinder sessions and/or they can be R1 devices. The migration controls evaluates the state of these devices to determine if the control operation can proceed.
- The devices in the storage group cannot be part of another migration session.

Geometry Compatible Mode flag set with Non-Disruptive Migration

Track size for FBA devices increased from 64K in Enginuity 5876 to 128K in HYPERMAX OS. Geometry Compatibility Mode (GCM) supports full SRDF functionality for devices on arrays running Enginuity 5876 with an odd number of cylinders paired with devices on arrays running HYPERMAX OS. With HYPERMAX OS, the device size difference is automatically managed using the device attribute Geometry Compatible Mode (GCM). A device with GCM set is presented as half a cylinder smaller than its true configured size, enabling full migration functionality between HYPERMAX OS and Enginuity 5876 for SRDF.

Because Non-Disruptive Migration uses SRDF technology for data migration (which requires that R1 and R2 devices be the same size), the GCM flag can get automatically set on the target device during migration.

Non-Disruptive Migration for PowerMaxOS 5978 overview

Non-Disruptive Migration (NDM) is a method for migrating data without application downtime. The migration takes place over a metro distance, typically within a data center.

Starting with PowerMaxOS there are two implementations of NDM each for different types of source array:

- VMAX3 or VMAX All Flash array running HYPERMAX OS 5977.1125.1125 or later with an ePack
- VMAX array running Enginuity 5876 with an ePack

When migrating to a PowerMax array, these are the only configurations for the source array.

Contact Dell EMC for the ePacks required for HYPERMAX OS 5977 and Enginuity 5876. In addition, the NDM support matrix has information on array operating systems support, host support, and multipathing support for NDM operations. The support matrix is available on the eLab Navigator.

Regulatory or business requirements for disaster recovery may require the use of replication to other arrays attached to source array, the target array, or both using SRDF/S, during the migration. In this case, contact Dell EMC for the ePacks required for the SRDF/S configuration.

Migration from a VMAX array

Migrating from a VMAX array uses SRDF technology. For NDM purposes, the source is a VMAX array running Enginuity 5876, with an ePack. The target is a PowerMax array running PowerMaxOS 5978. The following picture shows the logical structure of a migration from VMAX including the connections required:




Figure 13. Configuration of a VMAX migration

Process

The steps in the migration process are:

1. Set up the environment – configure the infrastructure of the source and target array, in preparation for data migration.
2. On the source array, select a storage group to migrate.
3. Create the migration session – copy the content of the storage group to the target array using SRDF.
4. Cutover the storage group to the PowerMax array.

From now on the application server accesses the storage group on the target array only.

5. Commit the migration session – remove resources from the source array and those used in the migration itself. The application now uses the target array only.
6. To migrate further storage groups, repeat steps 2 to 5.
7. After migrating all the required storage groups, remove the migration environment.

Other features

Other features of migrating from VMAX to PowerMax are:

- Data can be compressed during migration to the PowerMax array
- Allows for non-disruptive revert to the source array
- There can be up to 50 migration sessions in progress simultaneously
- NDM does not require an additional license as it is part of PowerMaxOS
- The connections between the application host and the arrays use FC; the SRDF connection between the arrays uses FC or GigE

Devices and components that cannot be part of a NDM system are:

- CKD devices, IBMi devices
- eNAS data
- ProtectPoint and FAST.X relationships along with their associated data

Migration from a VMAX3 or VMAX All Flash array

Migrating from VMAX3 or VMAX All Flash uses a modified form of SRDF/Metro. This means that in the normal workflow, both the source and target arrays are visible to the application host while the migration takes place. Indeed, both arrays are read/write accessible to the host. The following picture shows the logical structure of a migration from VMAX3 or VMAX All Flash including the connections required.

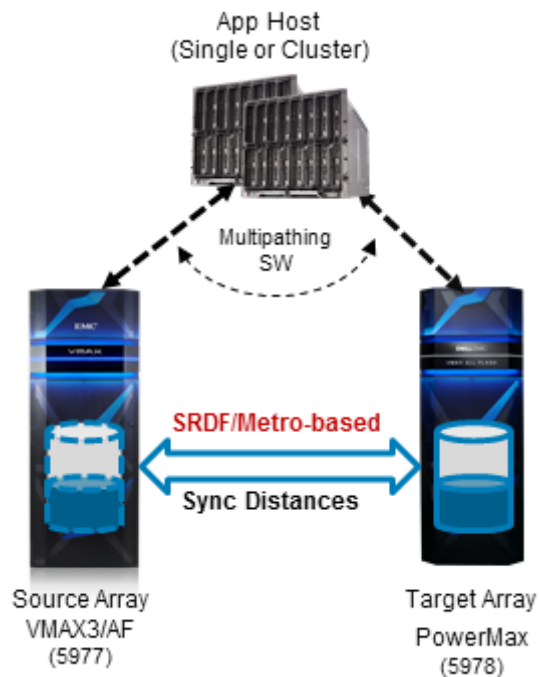


Figure 14. Configuration of a VMAX3 or VMAX All Flash migration

Process

Normal flow

The steps in the migration process that is normally followed are:

1. Set up the migration environment – configure the infrastructure of the source and target array, in preparation for data migration.
2. On the source array, select a storage group to migrate.

3. Create the migration session – copy the content of the storage group to the target array using SRDF/Metro.
During this time the source and target arrays are both accessible to the application host.
4. Commit the migration session – remove resources from the source array and those used in the migration itself. The application now uses the target array only.
5. To migrate further storage groups, repeat steps 2 to 4.
6. After migrating all the required storage groups, remove the migration environment.

Alternate flow

There is an alternative process that pre-copies the data to the target array before making it available to the application host. The steps in this process are:

1. Set up the migration environment – configure the infrastructure of the source and target array, in preparation for data migration.
2. On the source array, select a storage group to migrate.
3. Use the precopy facility of NDM to copy the selected data to the target array.
During this time the source array is available to the application host, but the target array is unavailable.
4. When the copying of the data is complete, use the Ready Target facility in NDM to make the target array available to the application host also.
5. Commit the migration session – remove resources from the source array and those used in the migration itself. The application now uses the target array only.
6. To migrate further storage groups, repeat steps 2 to 5.
7. After migrating all the required storage groups, remove the migration environment.

Other functions

Other NDM facilities that are available for exceptional circumstances are:

- Cancel – to cancel a migration that has not yet been committed.
- Sync – to stop or start the synchronization of writes to the target array back to source array. When stopped, the application runs on the target array only. Used for testing.
- Recover – to recover a migration process following an error.

Other features

Other features of migrating from VMAX3 or VMAX All Flash to PowerMax are:

- Data can be compressed during migration to the PowerMax array
- Allows for non-disruptive revert to the source array
- There can be up to 50 migration sessions in progress simultaneously
- Does not require an additional license as NDM is part of PowerMaxOS
- The connections between the application host and the arrays use FC; the SRDF connection between the arrays uses FC or GigE

Devices and components that cannot be part of a NDM system are:

- CKD devices, IBMi devices
- eNAS data
- ProtectPoint and FAST.X relationships along with their associated data

Environmental requirements for non-disruptive migration

There are requirements associated with both the arrays involved in a migration and the host system.

Storage arrays

- The target array runs PowerMaxOS 5978.
- The source array is either:
 - A VMAX3 or VMAX All Flash array running HYPERMAX OS 5977.1125.1125
 - A VMAX array running Enginuity 5876

The source array requires an ePack. Contact Dell EMC to obtain that ePack.

- SRDF is used for data migration, so zoning of SRDF ports between the source and target arrays is required. Note that an SRDF license is not required, as there is no charge for NDM.

- The NDM RDF group requires a minimum of two paths on different directors for redundancy and fault tolerance. If more paths are found up to eight paths are configured.
- If SRDF is not normally used in the migration environment, it may be necessary to install and configure RDF directors and ports on both the source and target arrays and physically configure SAN connectivity.

Host

- Wherever possible, use a host system separate from the application host to initiate and control the migration (the control host).
- The control host requires visibility of and access to both the source and target arrays.

Non-Disruptive Migration operational restrictions and state reference

This section details Non-Disruptive Migration states, valid migration states required for migration actions, and migration actions with other replication (TimeFinder and SRDF).

Non-Disruptive Migration session states

The following table lists the possible states for a migration session.

Migration state	Description
No migration	No migration in progress.
CreateInProg	Migration session is being created. The create command has not run to completion with either a success or a failed status.
Created	The create command has run to completion with a success status. [On VMAX to VMAX3 or VMAX AF]
CreateFailed	The create command has run to completion with a failed status.
CutoverReady	The create command has run to completion with a success status. A Host Discovery was performed, either automatically or manually. The devices on the target array are in pass-through mode and I/O's can be serviced by devices on the source and target arrays; a cutover may be performed. [On VMAX to VMAX3 or VMAX AF]
CutoverInProg	The cutover command is in progress and the application is being moved to the target array. The cutover command has not run to completion with either a success or a failed status. [On VMAX to VMAX3 or VMAX AF]
CutoverFailed	The cutover command has run to completion with a failed status. [On VMAX to VMAX3 or VMAX AF]
Precopy	The create command with the 'precopy' flag succeeded. The replication pathway has been created and is currently synchronizing data from the Source to the Target array. The Target devices are not masked to the host. A ReadyTgt may be performed. [On VMAX3 to VMAX AF].
ReadyTgtInProg	The ReadyTgt command is in progress and the application is being made host visible to the target array. The ReadyTgt command has not run to completion with either a success or a failed status. [On VMAX3 to VMAX AF].
ReadyTgtFailed	The ReadyTgt command has run to completion with a failed status. [On VMAX3 to VMAX AF].
Migrating	VMAX to VMAX3 or VMAX AF: The Cutover command succeeded and data is being migrated from the source to the target array. I/O's can only be serviced by devices on the target array. VMAX3 to VMAX AF: The Create or ReadyTgt command succeeded and data is being migrated from the source to the target array. I/O's can be serviced by devices on the source and target arrays.
MigrateFailed	This state occurs when a Migrating state is interrupted, as might occur with a loss of the required DM connectivity between the source and target arrays.
CutoverNoSync	The sync -stop command completed successfully. The application is running on the target array, all data was migrated from the source to the target array but data updates are not being replicated back to the source.

Migration state	Description
CutoverSync	The cutover or sync -start command completed successfully. The application is running on the target array, and all data has been synchronized between the target and source arrays.
CutoverSyncing	VMAX to VMAX3 or VMAX AF: The sync -start command completed successfully. The application is running on the target array, and data is being replicated back to the source. VMAX3 to VMAX AF: The sync -start command completed successfully. The application is running on both the source and target arrays and data is being synchronized between the source and target arrays.
Synchronized	The application is running on both the source and target arrays, a Host Discovery has being performed either automatically or manually and all data has been synchronized between the source and target arrays. [On VMAX3 to VMAX AF].
CommitInProg	Commit command is in progress. The commit command has not run to completion with either a success or a failed status.
CommitFailed	Commit command has run to completion with a failed status.
CancelInProg	Cancel command is in progress. The cancel command has not run to completion with either a success or a failed status.
CancelFailed	Cancel command has run to completion with a failed status.
Partitioned	The migration session has been successfully created but the DM replication pathway is not available.

Non-Disruptive Migration control actions and dependent migration states

Migration controls

Each migration action is dependent on the following and determines whether a migration action can proceed or fail:

- Migration state.
- Rules based on other replication states.
- Changes in the application storage configuration.
- Changes to the migration session outside the NDM feature.

Migration control actions and states

The following table lists the migration state that is a valid state prior to running a specific migration action. This table does not include the environment setup and remove actions as the `environment -setup` action must be completed successfully before any of the actions shown in the table can be performed. The `environment -remove` action removes the migration infrastructure, and all migration actions must be completed before performing the `environment -remove` action.

Table 25. Migration control actions and applicable migration states

	Migration states																						
	No migration	CreateInProg	Created ^a	CreateFailed	CutoverReady ¹	CutoverInProg ¹	CutoverFailed ¹	Precopy ²	ReadyTgtInProg ²	ReadyTgtFailed ²	Migrating	MigrateFailed	Synchronized ²	CutoverNoSync	CutoverSyncing	CutoverSync ¹	CommitInProg	CommitFailed	CancelInProg	CancelFailed	Partitioned		
create	Y																						
cancel			Y	Y	Y		Y				Y ^b	Y ^c	Y			Y ^c							
cutover					Y																		

Table 25. Migration control actions and applicable migration states (continued)

	Migration states																					
	No migration	CreateInProg	Created ^{1 a}	CreateFailed	CutoverReady ¹	CutoverInProg ¹	CutoverFailed ¹	Precopy ²	ReadyTgtInProg ²	ReadyTgtFailed ²	Migrating	MigrateFailed	Synchronized ²	CutoverNoSync	CutoverSyncing	CutoverSync ¹	CommitInProg	CommitFailed	CancelInProg	CancelFailed	Partitioned	
readytgt							Y															
sync -start														Y								
sync -stop													Y		Y	Y						
commit													Y	Y		Y						
recover				Y			Y		Y		Y						Y		Y			

- a. VMAX to VMAX3 or VMAX All Flash only.
- b. Revert flag required on VMAX to VMAX3 or VMAX All Flash only.
- c. VMAX3 to VMAX All Flash only.

Non-Disruptive Migration compatibility with other replication technologies

SRDF

SRDF relationships can exist on the migration source or target devices, however the following rules and restrictions apply.

- For the migration source device:
 - Can be a R1 device prior to the migration create action. After the create action, this device is seen as an R21 device, with the R2 device mirror used for the migration session.
 - Existing R1 device can be in Adaptive Copy, Asynchronous mode, or Synchronous mode prior to the migration create action.
 - RDF set mode action can be used to change to Adaptive Copy, Asynchronous mode, or Synchronous mode during the life cycle of the migration session.
 - Cannot be a R2 device prior to the create action, or at any time during the migration session.
 - Multi-session Consistency (MSC) cannot be enabled for the existing R1 device.
 - Existing R1 device cannot be part of a Star Consistency Group.
 - Data synchronization from the existing R2 to R1 device is not allowed at any time during the migration session.
 - Cannot add an R1 RDF mirror to the migration once it has started, therefore the NDM source device cannot be made an R21 device when the migration is in progress.
 - Existing R1 device cannot be enabled for Synchronous RDF Consistency (RDF-ECA).
 - Cannot change RDF mode from Asynchronous to Synchronous mode or Synchronous to Asynchronous on the existing R1 device.
- For the migration target device:
 - An additional R1 mirror can be added after the migration session is in the CutoverSync state. At this point in the migration session the device is seen as a R11 device with one mirror used for the migration session.
 - Added R1 device can be in Adaptive Copy, Asynchronous mode, or Synchronous mode.
 - Cannot add an additional R2 mirror (makes it a R21 device) at any time during the migration session.
 - RDF set mode action can be used to change to Adaptive Copy, Asynchronous mode, or Synchronous mode during the life cycle of the migration session.
 - Added R1 device cannot be part of an RDF Metro configuration.
 - Multi-session Consistency (MSC) cannot be enabled for the added R1 device.
 - Added R1 device cannot be part of a Star Consistency Group.
 - Data synchronization from the added R2 to R1 device is not allowed at any time during the migration session.
 - Added R1 device cannot be enabled for Synchronous RDF Consistency (RDF-ECA).
 - Cannot change RDF mode from Asynchronous to Synchronous mode or Synchronous to Asynchronous on the existing R1 device.

TimeFinder

TimeFinder relationships can exist on the migration source or target devices, however the following rules and restrictions apply.

- For the migration source device:
 - Can be a TimeFinder source device.
 - Cannot be a TimeFinder target device.
 - The TimeFinder session cannot be restoring data back to the migration source device at any time during the migration session.
- For the migration target device:
 - Can be a TimeFinder source device after the migration is in a CutoverSync state.
 - The TimeFinder session cannot be restored back to the migration target device at any time during the migration session.
 - The TimeFinder session on the target device must be removed prior to a cancel operation.

ORS

The source and target devices cannot be part of an ORS relationship.

Non-Disruptive Migration restrictions with other SYMCLI commands

Base commands

For the base commands, - symdev, symdg, symcg, symsg the following command options cannot be issued when the device is involved in a migration operation:

- write_disable
- not_ready
- free -all

For base commands listed above, the ready command option can be used to clear the host_inactive state which may have been placed on a device during a migration operation. This command should only be used to manually recover from a migration failure that cannot be resolved using the migration functionality. If the device is mapped to a host, the symforce flag is required.

NOTE: The ready command option only works on a device in the host_inactive state. For any other device state it will not clear the inactive state.

SRDF commands (symrdf)

The following actions are allowed for the symrdf set and control commands when an RDF pair is part of the a NDM RDF group:

- Setting link limbo
- Setting hardware compression
- Setting software compression
- Modifying a NDM RDF group

The following actions are not allowed for the symrdf set and control commands when an RDF pair is part of the a NDM RDF group:

- Creating RDF pairs in a NDM RDF group
- All RDF control actions
- All RDF set actions, except for link limbo
- Removing a NDM RDF group

TimeFinder commands (symmir, symclone, symsnapvx)

There are no changes to the syntax of these commands. However, these commands are enhanced to block all control actions that copy data to a NDM device, this includes not allowing the restore action if the NDM device is a source of a symmir, symclone, or symsnapvx session.

ORS commands (symrcopy)

There are no changes to the syntax of this command. However, this command is enhanced to block all control actions that copy data to a NDM device.

Non-Disruptive Migration control operations

The migration process is managed by the `symdm` CLI command and includes the following operations:

- Environment setup — Configures source and target array infrastructure for the migration process.
- Create — Replicates the application storage environment from source array to target array.
 - Create -precopy — (Only for HYPERMAX OS 5977 to PowerMaxOS) Starts copying the source data to the target array but does not make the target devices host visible. This allows the application to continue running only on the source array while allowing the data to be copied to the target. It requires a `readytgt` command to continue the migration.
- Readytgt — (Only for HYPERMAX OS 5977 to PowerMaxOS) Makes the target devices visible to the host, configures the DM replication so the I/Os are replicated to the other array and allows the I/Os to be serviced from both the source and target devices. This operation is only used after the `create -precopy` operation.
- Cutover — (Only for Enginuity 5876 to HYPERMAX OS 5977, or to PowerMaxOS) Switches the application data access form the source array to the target array and duplicates the application data on the source array to the target array.
- Commit — Removes application resources from the source array and releases the resources used for migration. Application permanently runs on the target array.

Migration sessions can also be reverted, recovered, or cancelled.

Environment setup for Non-Disruptive Migration

The environment configuration for NDM provides the following operations:

- Verify — Validates source and target array infrastructure for the migration process.
- Setup — Configures source and target array infrastructure (replication pathway) for the migration process.
- Remove — Once migration is committed or NDM environment no longer needed, removes the replication pathway configured for the migration.

Migration infrastructure - migration replication pathway creation

The environment setup operation creates an RDF (DM RDF) group to serve as a replication pathway between the source and target arrays. For Enginuity 5876 to HYPERMAX OS 5977, or to PowerMaxOS migrations this single replication group is used for moving the data between the source and target arrays. For HYPERMAX OS 5977 to PowerMaxOS migrations, this RDF group is used as a template; another RDF group is created for each migration to move the data between the source and target arrays.

The following configuration rules apply:

- The DM RDF group is configured with a minimum of two paths on different directors for redundancy and fault tolerance. If more paths are found up to eight paths will be configured.
- RF and RE ports are supported. If both are available the RF ports are selected.
- A single DM RDF group can be used for concurrent migrations between two arrays.
- For Enginuity 5876 to HYPERMAX OS 5977, or to PowerMaxOS migrations, changing or deleting the DM RDF group will cause any subsequent migration command to fail.
- A single target array can have multiple DM RDF groups, each connected to a different source array, with the target array setup as the migration target from multiple source arrays.
- A single source array can have multiple DM RDF groups each connected to a different target array, with the source array setup as the migration source to multiple target arrays.
- The source and target array must not be more than one hop away from the control host.

Validate environment for Non-Disruptive Migration

Description

The `symdm environment -validate` command validates that the environment is set up and meets the requirements for the migration process.

Syntax

To validate the environment for NDM, use the following syntax:

 **NOTE: The source and target array IDs must be specified.**

```
symdm environment -src_sid <SymID> -tgt_sid <SymID> -validate - noprompt
```

Options

-noprompt (-nop)

If included in command, user confirmation is required for command execution.

Examples

To verify the migration environment between source array 124 and target array 643, enter:

```
symdm environment -src_sid 124 -tgt_sid 643 -validate
```

```
A DM 'Validate Create' operation is  
in progress for storage group 'DM_APP1'. Please wait...
```

```
Analyze Configuration.....Validated.  
Initialize Replication Environment.....Validated.  
Create Storage Group(s) on Target.....Validated.  
Duplicate Device(s) on Target.....Validated.  
Create Initiator Group(s) on Target.....Validated.  
Create Port Group(s) on Target.....Failed.  
Create Masking View(s) on Target.....Validated.
```

```
The Validation failed. Please see the SYMAPI log file for more information.
```

Expected response when running environment -validate command on an environment that is not setup for migration:

```
symdm environment -src_sid 124 -tgt_sid 643 -validate
```

```
A DM 'Environment Validate' operation is  
in progress. Please wait...
```

```
Analyze Configuration.....Failed.
```

```
The Validation failed. Please see the SYMAPI log file for more information.
```

Setup environment for Non-Disruptive Migration

Description

The `symdm environment -setup` command configures the source and target array infrastructure (replication pathway) for the migration process.

Syntax

To configure the environment for NDM, use the following syntax:

```
symdm environment -src_sid <SymID> -tgt_sid <SymID> -setup
```

Examples

To setup the migration path between source array 124 and target array 643, enter:

```
symdm environment -src_sid 124 -tgt_sid 643 -setup
```

```
A DM 'Environment Setup' operation is  
in progress. Please wait...
```

```
Analyze Configuration.....Started.
```

```
Source SID:000198700124
Target SID:000197100643
Analyze Configuration.....Done.
Setup Configuration.....Started.
Setup Configuration.....Done.
```

Expected response when running `environment-setup` command and the migration environment is already setup:

```
symdm environment -src_sid 124 -tgt_sid 643 -setup


A DM 'Environment Setup' operation is in progress. Please wait...

The migration environment is already configured.
```

Migration session creation for Non-Disruptive Migration

The migration session creation for NDM provides the following operations:

- **Validate** — Checks that the requirements and restrictions are met for the source array and the target array to ensure that a specific migration (storage group) can proceed. This command does not actually start the migration process.
- **Create** — performs several actions:
 - Provisions the target array with the equivalent storage used on the source array for the application.
 - Configures the target array devices to have the same identity as the source array devices.
 - Duplicates the application storage environment, of the specified storage group, on the source array to the target array.

 **NOTE:** Once the migration operation starts, do not reboot hosts or modify storage groups.

Validate create session for Non-Disruptive Migration

Description

The `symdm create -validate` command validates that a specific migration (storage group) can proceed and replicates the application storage environment, of the specified storage group, on the source array to the target array.

Syntax

To validate the create action for NDM, use the following syntax:

```
symdm create -src_sid <SymID> -tgt_sid <SymID> -sg<sgName> -validate -noprompt
```

Options

-noprompt (-nop)

If included in command, user confirmation is not required for command execution.

Examples

To validate migration of storage group `DM_APP1` between source array 124 and target array 643, enter:

```
symdm create -src_sid 124 -tgt_sid 643 -sg DM_APP1 -validate

A DM 'Validate Create' operation is
in progress for storage group 'DM_APP1'. Please wait...

Analyze Configuration.....Validated.
Create Storage Group(s) on Target.....Validated.
Duplicate Device(s) on Target.....Validated.
Create Initiator Group(s) on Target.....Validated.
Create Port Group(s) on Target.....Validated.
Create Masking View(s) on Target.....Validated

The DM 'Validate Create' operation successfully executed for
storage group 'DM_APP1'
```

Expected response when running `create -validate` that fails:

```
symdm create -src_sid 124 -tgt_sid 643 -sg DM_APP1 -validate

A DM 'Validate Create' operation is
in progress for storage group 'DM_APP1'. Please wait...

  Analyze Configuration.....Validated.
  Create Storage Group(s) on Target.....Validated.
  Duplicate Device(s) on Target.....Validated.
  Create Initiator Group(s) on Target.....Validated.
  Create Port Group(s) on Target.....Failed.
  Create Masking View(s) on Target.....Validated

The Validation failed. Please see the SYMAPI log file for more information.
```

The above `create -validate` command indicates a problem with the port-create action. This must be corrected before the migration is allowed.

Create session for Non-Disruptive Migration

Description

The `symdm create` command duplicates the application data, of the specified storage group, on the source array to the target array.

Syntax

To create a NDM session, use the following syntax:

```
symdm create -src_sid <SymID> -tgt_sid <SymID> -sg <sgName> -noprompt
-tgt_srp <SRPName> -tgt_pg <PgName> -nocompression -precopy
```

Options

`-noprompt (-nop)`

If included in command, user confirmation is not required for command execution.

`-tgt_srp`

Specifies the SRP name to use on the target array during the NDM create operation.

`-tgt_pg`

Specifies the PG name to use on the target array during the NDM create operation.

`-nocompression`

If included in command, the compression attribute is removed.

`-precopy`

If included in command, NDM starts copying the source data to the target array but does not make the target devices host visible. This is only available for HYPERMAX OS 5977 to PowerMaxOS migration.

Examples

To create a NDM session between source array 124 and target array 643 for storage group DM-APP1, enter:

```
symdm create -src_sid 124 -tgt_sid 643 -sg DM_APP1

A DM 'Create' operation is in progress for storage group 'DM_APP1'. Please wait...

  Analyze Configuration.....Started.
    Source SID:000198700124
    Target SID:000197100643
  Analyze Configuration.....Done.
  Set Dynamic RDF attribute on Source Device(s).....Started.
  Set Dynamic RDF attribute on Source Device(s).....Done.
  Create Storage Group(s) on Target.....Started.
  Create Storage Group(s) on Target.....Done.
  Duplicate Device(s) on Target.....Started.
  Duplicate Device(s) on Target.....In Progress.
```



```

Duplicate Device(s) on Target.....Done.
Create Initiator Group(s) on Target.....Started.
Create Initiator Group(s) on Target.....Done.
Create Port Group(s) on Target.....Started.
Create Port Group(s) on Target.....Done.
Setup Data Replication.....Started.
Setup Data Replication.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.
Update Device State.....Started.
Update Device State.....Done.

```

The DM 'Create' operation successfully executed for storage group 'DM_APP1'

Expected response when issuing a create command for a migration that already in progress:

```
symdm create -src_sid 124 -tgt_sid 643 -sg DM_APP1
```

A DM 'Create' operation is in progress for storage group 'DM_APP1'. Please wait...

The migration session is already in the requested state.

Expected response when issuing a create command for a migration and the storage group is not in a masking view:

```
symdm create -src_sid 124 -tgt_sid 643 -sg DM_APP1
```

A DM 'Create' operation is in progress for storage group 'DM_APP1'. Please wait...

```

Analyze Configuration.....Started.
Source SID:000198700124
Target SID:000197100643
Analyze Configuration.....Failed.


```

The storage group is not in a masking view

Create precopy session for Non-Disruptive Migration (only for HYPERMAX OS 5977 to PowerMaxOS migration)

Description

The `symdm create -precopy` command configures the replication to start copying the source data to the target array after provisioning storage on the target array but does not make the target devices host visible. This allows the application to continue running only on the source array while allowing the data to be copied to the target.

 **NOTE:** If the `-precopy` option is used, then the `readytgt` command is required to continue the migration process.

Syntax

To create a precopy session, use the following syntax:

```
symdm create -precopy -src_sid <SymID> -tgt_sid <SymID> -sg <sgName> -noprompt
```

Options

`-noprompt (-nop)`

If included in command, user confirmation is not required for command execution.

Examples

To create a precopy session between source array 124 and target array 643 for storage group DM-APP1, enter:

```
symdm create -precopy -src_sid 124 -tgt_sid 643 -sg DM_APP1
```

A DM 'Precopy Create' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
Analyze Configuration.....Done.
Initialize Replication Environment.....Started.
Initialize Replication Environment.....Done.
Create Storage Group(s) on Target.....Started.
Create Storage Group(s) on Target.....Done.
Duplicate Device(s) on Target.....Started.
Preparing for device create on Target.....Started.
Preparing for device create on Target.....Done.
Duplicate Device(s) on Target.....In Progress.
Duplicate Device(s) on Target.....Done.
Create Initiator Group(s) on Target.....Started.
Create Initiator Group(s) on Target.....Done.
Create Port Group(s) on Target.....Started.
Create Port Group(s) on Target.....Done.
Start Data Replication.....Started.
Start Data Replication.....Done.
```

The DM 'Precopy Create' operation successfully executed for storage group 'DM_APP1'.

Cutover Non-Disruptive Migration session (only for Enginuity 5876 to HYPERMAX OS 5977, or to PowerMaxOS migration)

Description

The `symdm cutover` command switches the application data access from the source array to the target array. After a successful cutover, the application will be running against newly configured storage on the target array, with updates to its data replicating back to the source array.

NOTE: Prior to running the cutover action, a host rescan must be performed to verify that new paths to the application's devices are discovered.

Syntax

To cutover a migration session, use the following syntax:

NOTE: The source or target array ID and the migrated storage group must be specified.

```
symdm cutover -sid <SymID> -sg <sgName> -noprompt
```

Options

`-noprompt`

If included in command, user confirmation is required for command execution.

Examples

To cutover the migration session for storage group DM_APP1 on target array 643, enter:

```
symdm cutover -sid 643 -sg DM_APP1
```

A DM 'Cutover' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000198700124
  Target SID:000197100643
Analyze Configuration.....Done.
Cutover.....Started.
```

```
Cutover.....Done.
```

```
The DM 'Cutover' operation successfully executed for storage group 'DM_APP1'
```

Expected response when issuing a cutover command for a migration that is in process:

```
symdm cutover sid 643 -sg DM_APP1
```

```
A DM 'Cutover' operation is in progress for storage group 'DM_APP1'. Please wait...
```

```
The migration session is already in the requested state.
```

Readytgt Non-Disruptive Migration session (only for HYPERMAX OS 5977 to PowerMaxOS migration)

Description

The `symdm readytgt` command reconfigures the DM replication so the I/Os are replicated to the other array and makes the target devices visible to the host allowing the I/Os to be serviced from both the source and target devices.

NOTE: This operation is only used if the `-precopy` option was used with the `create` operation. After successfully completing the `readytgt` operation, a host rescan must be performed to ensure that the application host can use the target-side devices for I/O.

At this point both the target-side and source-side devices are `host_active`, I/O's can be serviced by either array and are replicated to the other array.

Syntax

To reconfigure the replication, use the following syntax:

NOTE: The source or target array ID and the migrated storage group must be specified.

```
symdm readytgt -sid <SymID> -sg <sgName> -i <Interval> -c <Count> -noprompt
```

Options

`-c <#>, -i <#>`

Executes list command for the specified number of times and the specified interval.

`-noprompt`

If included in command, user confirmation is not required for command execution.

Examples

To reconfigure the migration session for storage group `DM_APP1` on target array 643, enter:

```
symdm readytgt -sid 643 -sg DM_APP1
A DM 'ReadyTgt' operation is
in progress for storage group 'DM_APP1'. Please wait...

Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
Analyze Configuration.....Done.
Preparing Devices for Host discovery.....Started.
Preparing Devices for Host discovery.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.
```

```
The DM 'ReadyTgt' operation successfully executed for
storage group 'DM_APP1'.
```

Expected response when issuing a readytgt command for a migration that is in process:

```
symdm readytg sid 643 -sg DM_APP1
```

```
A DM 'ReadyTgt' operation is in progress for storage group 'DM_APP1'. Please wait...
```

```
The migration session is already in the requested state.
```

Commit Non-Disruptive Migration session

Description

The `symdm commit` command completes the migration by removing application resources from the source array and releasing resources used for the migration.

NOTE: After the commit operation is complete, a host rescan that includes the `-r` option should be run to remove the dead paths to the target array, This should be followed by a multipath reload command.

Syntax

To commit a migration, use the following syntax:

NOTE: The source or target array ID and the migrated storage group must be specified.

```
symdm commit -sid <SymID> -sg <sgName> -noprompt
```

Options

`-noprompt`

If included in command, user confirmation is required for command execution.

Examples

To commit the migration for storage group `DM_APP1` on target array 643, enter:

```
symdm commit -sid 643 -sg DM_APP1
```

```
A DM 'Commit' operation is in progress for storage group 'DM_APP1'. Please wait...
```

```
Analyze Configuration.....Started.  
  Source SID:000198700124  
  Target SID:000197100643  
Analyze Configuration.....Done.  
Remove Masking View(s) on Source.....Started.  
Remove Masking View(s) on Source.....Done.  
Remove Data Replication.....Started.  
Remove Data Replication.....In Progress.  
Remove Data Replication.....Done.
```

```
The DM 'Commit' operation successfully executed for storage group 'DM_APP1'
```

If a commit operation fails, correct the cause of the failure and run a recover operation.

Post-commit infrastructure actions

After a commit operation is performed, following infrastructure changes take place:

- RDF device relationships used for migrating data via the DM RDF group are deleted, however the DM RDF group remains in place.
- Masking views on the source array are removed if there are no dedicated gatekeepers in the storage group.
- Source-side devices are assigned the device ID of the target-side device to ensure that the device is no longer used by the application that was moved to the target.
- If a migrated storage group had dedicated gatekeepers, a new storage group with the name of the migrated storage group and a suffix `_SAVE_#` is created on the source array, and all migrated devices on the source array are moved to this storage group.

Cancel Non-Disruptive Migration session

Description

The `symdm cancel` command halts a migration that has not been committed. The cancel operation applies only to the migration of the specified storage group. Other migrations running in parallel will continue, and new migrations can be created at any time during or after the cancel operation. A cancel action is blocked if the target-side devices are configured for local or remote replication. Replication must be removed prior to running the cancel operation.

NOTE: After the cancel operation is complete, a host rescan that includes the option to remove the dead paths to the target array. This should be followed by a multipath reload command.

For Engenuity 5876 to HYPERMAX OS 5977, or to PowerMaxOS migrations, if a cutover command has been issued, the cancel action must use the `-revert` flag. This forces all data to be replicated back to the source side devices before the cancel can complete.

Syntax

To cancel a migration, use the following syntax:

NOTE: The source or target array ID and the migrated storage group must be specified.

```
symdm cancel -sid <SymID> -sg <sgName> -noprompt
```

Options

`-noprompt`

If included in command, user confirmation is not required for command execution.

Examples

To cancel the migration for storage group `DM_APP1` on target array 643 prior to cutover, enter:

```
symdm cancel -sid 643 -sg DM_APP1
```

```
A DM 'Cancel' operation is in progress for storage group 'DM_APP1'. Please wait...
```

```
Analyze Configuration.....Started.
  Source SID:000198700124
  Target SID:000197100643
Analyze Configuration.....Done.
Remove Masking View(s) on Target.....Started.
Remove Masking View(s) on Target.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Port Group(s) on Target.....Started.
Remove Port Group(s) on Target.....Done.
Remove Initiator Group(s) on Target.....Started.
Remove Initiator Group(s) on Target.....Done.
Remove Duplicate Device(s) on Target.....Started.
Wait for deallocation to complete .....Started.
Wait for deallocation to complete .....Done.
Remove Duplicate Device(s) on Target.....Done.
Remove Storage Group(s) on Target.....Started.
Remove Storage Group(s) on Target.....Done.
```

```
The DM 'Cancel' operation successfully executed for storage group 'DM_APP1'
```

To cancel the migration, that has been cutover, for storage group `DM_APP1` on target array 643, enter:

```
symdm cancel -revert -sid 643 -sg DM_APP1
```

```
A DM 'Cancel Revert' operation is
in progress for storage group 'DM_APP1'. Please wait...
```

```
Analyze Configuration.....Started.
  Source SID:000198700124
```

```

Target SID:000197100643
Analyze Configuration.....Done.
Revert Data Replication.....Started.
Revert Data Replication.....In Progress.
Revert Data Replication.....Done.
Remove Masking View(s) on Target.....Started.
.....

```

The DM 'Cancel Revert' operation successfully executed for storage group 'DM_APP1'

Upon successful completion of a cancel operation, the storage environment reverts to the pre-migration state:

- Migration replication pathway connections are severed.
- Any target-side resources configured for the application that are not shared with other applications are removed.
- Application is running on the source only.

Synchronizing devices for Non-Disruptive Migration session

The sync operation controls target to source replication after all of the data is on the target array.

If the target-to-source device replication needs to be controlled, the following sync control actions are available:

- `sync -stop` — Suspends target-to-source device replication and puts the migration in a *CutoverNoSync* state.
- `sync -start` — Establishes target-to-source device replication and puts the migration in a *CutoverSyncing* or *CutoverSync* state.

Stop device synchronization for Non-Disruptive Migration session

Description

The `symdm sync -stop` command suspends target-to-source replication.

Syntax

To stop target-to-source replication, use the following syntax:

 **NOTE:** The source or target array ID, and the migrating storage group must be specified.

```
symdm sync -sid <SymID> -sg <sgName> -stop
```

Examples

To stop target-to-source replication for storage group DM-APP1 from target array 643, enter:

```
symdm sync -sid 643 -sg DM_APP1 -stop
```

A DM 'Sync Stop' operation is in progress for storage group 'DM_APP1'. Please wait...

```

Analyze Configuration.....Started.
Source SID:000198700124
Target SID:000197100643
Analyze Configuration.....Done.
Stop Data Replication.....Started.
Stop Data Replication.....Done.

```

The DM 'Sync Stop' operation successfully executed for storage group 'DM_APP1'

Start device synchronization for Non-Disruptive Migration session

Description

The `symdm sync -start` command establishes target-to-source replication.

Syntax

To start target-to-source replication, use the following syntax:

NOTE: The source or target array ID, and the migrating storage group must be specified.

```
symdm sync -sid <SymID> -sg <sgName> -start
```

Examples

To start target-to-source replication for storage group DM-APP1 from target array 643, enter:

```
symdm sync -sid 643 -sg DM_APP1 -start
```

```
A DM 'Sync Start' operation is in progress for storage group 'DM_APP1'. Please wait...
```

```
Analyze Configuration.....Started.  
  Source SID:000198700124  
  Target SID:000197100643  
Analyze Configuration.....Done.  
Start Data Replication.....Started.  
Start Data Replication.....Done.
```

```
The DM 'Sync Start' operation successfully executed for storage group 'DM_APP1'
```

Recover a failed Non-Disruptive Migration session

Description

The `symdm recover` command recovers the last action performed and puts the migration in a state that allows the failed action (create, cutover, commit, or cancel) to complete. The recover operation is only used after a migration step results in a *Failed* state.

A recover operation performs the following actions:

- Determines which migration step failed.
- Puts the migration session's resources (connections, devices, etc.) into the appropriate state to allow the failed action to complete.
- Repeats or resumes (depending on the cause of the failure) the failed action.

Syntax

To recover from a failed migration, use the following syntax:

NOTE: The source or target array ID and the migrated storage group must be specified.

```
symdm recover -sid <SymID> -sg <sgName>
```

Examples

To recover from a failed create migration of storage group DM_APP1 on target array 643, enter:

```
symdm recover -sid 643 -sg DM_APP1
```

```
A DM 'Recover' operation is in progress for storage group 'DM_APP1'. Please wait...
```

```
Analyze Configuration.....Started.  
  Source SID:000198700124  
  Target SID:000197100643  
Analyze Configuration.....Done.  
Create Storage Group(s) on Target.....Not Needed.  
Duplicate Device(s) on Target.....Not Needed.  
Create Initiator Group(s) on Target.....Not Needed.  
Create Port Group(s) on Target.....Not Needed.  
Setup Data Replication.....Started.  
Setup Data Replication.....Done.
```

```
Recover Data Replication.....Started.
Recover Data Replication.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.
Update Device State.....Started.
Update Device State.....Done.
```

The DM 'Recover' operation successfully executed for storage group 'DM_APP1'

To recover from a failed cutover of storage group DM_APP1 on target array 643, enter:

```
symdm recover -sid 643 -sg DM_APP1
```

A DM 'Recover' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000198700124
  Target SID:000197100643
Analyze Configuration.....Done.
Recover Data Replication.....Started.
Recover Data Replication.....Done.
Cutover.....Started.
Cutover.....In Progress.
Cutover.....In Progress.
Cutover.....Done.
```

The DM 'Recover' operation successfully executed for storage group 'DM_APP1'

To recover from a failed commit of storage group DM_APP1 on target array 643, enter:

```
symdm recover -sid 643 -sg DM_APP1
```

A DM 'Recover' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000198700124
  Target SID:000197100643
Analyze Configuration.....Done.
Remove Masking View(s) on Source.....Started.
Remove Masking View(s) on Source.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
```

The DM 'Recover' operation successfully executed for storage group 'DM_APP1'

To recover from a failed cancel for the migration session of storage group DM_APP1 on target array 643, enter:

```
symdm recover -sid 643 -sg DM_APP1
```

A DM 'Recover' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000198700124
  Target SID:000197100643
Analyze Configuration.....Done.
Remove Masking View(s) on Target.....Not Needed.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Port Group(s) on Target.....Started.
Remove Port Group(s) on Target.....Done.
Remove Initiator Group(s) on Target.....Started.
Remove Initiator Group(s) on Target.....Done.
Remove Duplicate Device(s) on Target.....Started.
Wait for deallocation to complete.....Started.
Wait for deallocation to complete.....Done.
Remove Duplicate Device(s) on Target.....Done.
Remove Storage Group(s) on Target.....Started.
Remove Storage Group(s) on Target.....Done.
```

The DM 'Recover' operation successfully executed for storage group 'DM_APP1'

To recover from a failed revert for the migration session of storage group DM_APP1 on target array 643, enter:

```
symdm recover -sid 643 -sg DM_APP1
```

A DM 'Recover' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000198700124
  Target SID:000197100643
Analyze Configuration.....Done.
Revert Data Replication.....Started.
Revert Data Replication.....In Progress.
Revert Data Replication.....Done.
Remove Masking View(s) on Target.....Started.
Remove Masking View(s) on Target.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Port Group(s) on Target.....Started.
Remove Port Group(s) on Target.....Done.
Remove Initiator Group(s) on Target.....Started.
Remove Initiator Group(s) on Target.....Done.
Remove Duplicate Device(s) on Target.....Started.
Wait for deallocation to complete. ....Started.
Wait for deallocation to complete. ....23181
Wait for deallocation to complete. ....0
Wait for deallocation to complete. ....0
Wait for deallocation to complete. ....Done.
Remove Duplicate Device(s) on Target.....Done.
Remove Storage Group(s) on Target.....Started.
Remove Storage Group(s) on Target.....Done.
```

The DM 'Recover' operation successfully executed for storage group 'DM_APP1'

Remove environment for Non-Disruptive Migration session

Description

The `symdm environment -remove` command removes the migration infrastructure created by the `-setup` option after all necessary application migrations have been completed. The migration must be cancelled or committed before this command is allowed.

NOTE: Separate remove operations are required for each source and target array pair. The environment remove operation removes the DM RDF group between the two arrays, provided there are no migration sessions in progress.

Syntax

To remove the environment for NDM, use the following syntax:

NOTE: The source and target array IDs must be specified.

```
symdm environment -src_sid <SymID> -tgt_sid <SymID> -remove
```

Examples

To remove the migration path between source array 124 and target array 643, enter:

```
symdm environment -src_sid 124 -tgt_sid 643 -remove
```

A DM 'Environment Remove' operation is in progress. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000198700124
  Target SID:000197100643
Analyze Configuration.....Done.
Remove Configuration.....Started.
Remove Configuration.....Done.
```

The DM 'Environment Remove' operation successfully executed.

List Non-Disruptive Migration session status

Syntax

To monitor migration session status, use the following syntax:

```
symdm list
```

Options

-sg <SgName>

Lists migration status for specified storage group only.

-c <#>, -i <#>

Executes list command for the specified number of times and the specified interval.

-v

Lists summary information about each of the current migration sessions and will display where the migration session failed.

-detail

Lists detailed information on the state of a single migration session. Displays a detailed list of all objects used in the migration session of a specified storage group (-sg). Using these options can help to isolate the cause of a failed migration.

-environment [-offline]

Reports summary information for configured migration environments. All local and remote arrays are queried. The -offline option sets this operation to run in offline mode using only the host configuration database.

-sg_info | -pg_info | -ig_info | -view_info | -pairs_info

Filters the `symdm list -sg -v - detail` command to list the detail for only storage groups, port groups, initiator groups, masking views, or device pairs.

Examples

To list the migrations sessions and the status, for array 084, enter:

```
symdm list -sid 084
```

To list the migration session and status for storage group `storgrp_a`, enter:

```
symdm list -sid 084 -sg storgrp_a
```

To list the summary information for each of the currently configured migration environments, enter:

```
symdm list -environment
```

Sample output

For all sessions in a specified array:

```
Symmetrix ID: 000197100084
```

Storage Group	Source Array	Target Array	State	Total Capacity (GB)	Done (%)
storgrp_a	000198700084	000197100085	Migrating	500.0	90

```
storgrp_b          000198700084 000197100085 CutoverFailed      1000.0  N/A
```

For specific storage group:

```
Symmetrix ID : 000197100084
Storage Group: storgrp_a
```

Source Symmetrix	Target Symmetrix	State	Total Capacity (GB)	Done (%)
000197100084	000197100085	Migrating	500.0	90

For migration environment status:

```
Symmetrix ID: 000194901137
```

Remote SymmID	Status
000196801476	OK
000197100084	Failed


```
Symmetrix ID: 000196801476
```

Remote SymmID	Status
000194901137	OK
000198700030	Failed

Using list command to troubleshoot migration failure

Refer to [Example: Troubleshoot migration failure \(device moved between storage groups\)](#) on page 411 and [Example: Troubleshoot migration failure \(masking view added\)](#) on page 413 for using the `symdm list` command for troubleshooting a migration failure.

Non-Disruptive Migration examples

This section provides the following examples for using Non-Disruptive Migration to migrate an application to a newly acquired VMAX All Flash array (ID 000197300015)

- [Example: Non-Disruptive Migration environment setup](#) on page 403
- [Example: Typical Non-Disruptive Migration process \(from Enginuity 5876 to HYPERMAX OS 5977, or to PowerMaxOS\)](#) on page 404
- [Example: Typical Non-Disruptive Migration process \(from HYPERMAX OS 5977 to PowerMaxOS 5978\)](#) on page 406
- [Example: Suspending, restarting and recovering the Non-Disruptive Migration session](#) on page 407
- [Example: Cancelling the Non-Disruptive Migration session](#) on page 410
- [Example: Troubleshoot migration failure \(device moved between storage groups\)](#) on page 411
- [Example: Troubleshoot migration failure \(masking view added\)](#) on page 413

Example: Non-Disruptive Migration environment setup

The following example shows how to use the `symdm environment -setup` command to prepare to migrate two applications each from a different source VMAX array to the same target VMAX3 array:

- HR2 currently uses VMAX array 000194902222 for data storage.
- PAYROLL4 currently uses VMAX array 000194904444 for data storage.
- Both applications will be migrated to VMAX All Flash array 000197300015.

Both migrations can be performed concurrently, but the migration infrastructure requires that concurrent migration setup and create operations are constructed separately. Therefore, the environment setup commands are run separately.

NOTE: When either environment setup completes, the migration for the completed setup can be started.

The environment setup operation is run for source array 222 and target array 015.

```
syndm environment -src_sid 222 -tgt_sid 015 -setup
```

```
A DM 'Environment Setup' operation is in progress. Please wait...
```

```
Analyze Configuration.....Started.  
  Source SID:000194902222  
  Target SID:000197300015  
Analyze Configuration.....Done.  
Setup Configuration.....Started.  
Setup Configuration.....Done.
```

```
The DM 'Environment Setup' operation successfully executed.
```

The environment setup operation is run for source array 444 and target array 015.

```
syndm environment -src_sid 444 -tgt_sid 015 -setup
```

```
A DM 'Environment Setup' operation is in progress. Please wait...
```

```
Analyze Configuration.....Started.  
  Source SID:000194904444  
  Target SID:000197300015  
Analyze Configuration.....Done.  
Setup Configuration.....Started.  
Setup Configuration.....Done.
```

```
The DM 'Environment Setup' operation successfully executed.
```

Example: Typical Non-Disruptive Migration process (from Enginuity 5876 to HYPERMAX OS 5977, or to PowerMaxOS)

The HR2 application is storing data on the source VMAX array 000194902222. The following example shows how to use the `syndm create`, `syndm cutover` and `syndm commit` commands to move the HR2 data target VMAX3 array 000197300015 and run the application from this array.

1. The create operation is run specifying the source array 222, target array 015, and the storage group HR2 as application data to be migrated.

```
syndm create -src_sid 222 -tgt_sid 015 -sg HR2.
```

```
A DM 'Create' operation is in progress for storage group 'HR2'. Please wait...
```

```
Analyze Configuration.....Started.  
  Source SID:000194902222  
  Target SID:000197300015  
Analyze Configuration.....Done.  
Create Storage Group(s) on Target.....Started.  
Create Storage Group(s) on Target.....Done.  
Duplicate Device(s) on Target.....Started.  
Duplicate Device(s) on Target.....In Progress.  
Duplicate Device(s) on Target.....Done.  
Create Initiator Group(s) on Target.....Started.  
Create Initiator Group(s) on Target.....Done.  
Create Port Group(s) on Target.....Started.  
Create Port Group(s) on Target.....Done.  
Setup Data Replication.....Started.  
Setup Data Replication.....Done.  
Create Masking View(s) on Target.....Started.  
Create Masking View(s) on Target.....Done.  
Update Device State.....Started.  
Update Device State.....Done.
```

```
The DM 'Create' operation successfully executed for storage group 'HR2'
```

2. A host rescan is run. After the host rescan:
 - The application can issue I/Os to both the source and target devices.
 - The host multi-pathing software discovers that HR2's storage accessible using a path to the target array.
 - The migration session is in CutoverReady state which indicates that the array has a minimum of one path from an initiator in a masking view.
3. A cutover operation is run.

NOTE: Either the source or the target array participating in the migration session can be specified in the `syndm cutover` command.

During the cutover operation the NDM process:

- Moves HR2 I/O processing to only the target array devices.
- Makes source array devices *host_inactive*.
- The host multi-pathing software discovers that HR2's storage is not accessible using the path to the source array.
- Writes issued to HR2's storage on the target array replicate back to the source array, in the event that the migration is cancelled and the HR2 processing reverts to the source array.

```
syndm cutover -sid 015 -sg HR2
```

```
A DM 'Cutover' operation is in progress for storage group 'HR2'. Please wait...
```

```
Analyze Configuration.....Started.
  Source SID:000194902222
  Target SID:000197300015
Analyze Configuration.....Done.
Cutover.....Started.
Cutover.....Done.
```

```
The DM 'Cutover' operation successfully executed for storage group 'HR2'
```

HR2 application is now running on the target array.

4. A list operation is run to check for the CutoverSync state for the migration session.

NOTE: Either the source or the target array participating in the migration session can be specified in the `syndm list` command.

```
syndm list -sid 222 -sg HR2
```

```
Symmetrix ID : 000194902222
Storage Group: storgrp_a
```

Source Array	Target Array	State	Total Capacity (GB)	Done (%)
000194902222	000197300015	Migrating	500.0	90

All HR2 data has been synchronized between the source and target arrays.

5. A commit operation is run to complete the migration of HR2.

NOTE: Either the source or the target array participating in the migration session can be specified in the `syndm commit` command.

```
syndm commit -sid 015 -sg HR2
```

```
A DM 'Commit' operation is in progress for storage group 'HR2'. Please wait...
```

```
Analyze Configuration.....Started.
  Source SID:000194902222
  Target SID:000197300015
Analyze Configuration.....Done.
Remove Masking View(s) on Source.....Started.
Remove Masking View(s) on Source.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....In Progress.
Remove Data Replication.....Done.
```

```
The DM 'Commit' operation successfully executed for storage group 'HR2'
```

HR2 application now permanently runs on the target array and no longer runs on the source array.

i **NOTE: Run a host rescan using `-r` option, followed by a multipath reload command.**

Example: Typical Non-Disruptive Migration process (from HYPERMAX OS 5977 to PowerMaxOS 5978)

The DM_APP1 application is storing data on the source VMAX array 000194902643. The following example shows how to use the `syndm create`, `syndm cutover` and `syndm commit` commands to move the DM_APP1 data target VMAX3 array 000197300644 and run the application from this array.

1. The `create -validate` operation is run to check if the requirements and restrictions are met to ensure that the specified migration can proceed.

```
syndm create -src_sid 643 -tgt_sid 644 -sg DM_APP1 -validate

A DM 'Validate Create' operation is
in progress for storage group 'DM_APP1'. Please wait...

Analyze Configuration.....Validated.
Initialize Replication Environment.....Validated.
Create Storage Group(s) on Target.....Validated.
Duplicate Device(s) on Target.....Validated.
Create Initiator Group(s) on Target.....Validated.
Create Port Group(s) on Target.....Validated.
Create Masking View(s) on Target.....Validated.
```

2. The `create` operation is run specifying the source array 643, target array 644, and the storage group DM_APP1 as application data to be migrated.

```
syndm create -src_sid 643 -tgt_sid 644 -sg DM_APP1

A DM 'Create' operation is
in progress for storage group 'DM_APP1'. Please wait...

Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
Analyze Configuration.....Done.
Initialize Replication Environment.....Started.
Initialize Replication Environment.....Done.
Create Storage Group(s) on Target.....Started.
Create Storage Group(s) on Target.....Done.
Duplicate Device(s) on Target.....Started.
Duplicate Device(s) on Target.....In Progress.
Duplicate Device(s) on Target.....Done.
Create Initiator Group(s) on Target.....Started.
Create Initiator Group(s) on Target.....Done.
Create Port Group(s) on Target.....Started.
Create Port Group(s) on Target.....Done.
Start Data Replication.....Started.
Start Data Replication.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.

The DM 'Create' operation successfully executed for
storage group 'DM_APP1'.
```

3. **OPTIONAL Precopy:** Using the `-precopy` option duplicates the DM_APP1 application's storage on the target array 644 and configures data replication from the source array 643 to the target array.

i **NOTE: The target array devices are not made host visible when the command completes. After the precopy operation, you**

```
syndm create -precopy -src_sid 643 -tgt_sid 644 -sg DM_APP1

A DM 'Precopy Create' operation is
in progress for storage group 'DM_APP1'. Please wait...

Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
```

```
Analyze Configuration.....Done.
Initialize Replication Environment.....Started.
Initialize Replication Environment.....Done.
Create Storage Group(s) on Target.....Started.
Create Storage Group(s) on Target.....Done.
Duplicate Device(s) on Target.....Started.
Preparing for device create on Target.....Started.
Preparing for device create on Target.....Done.
Duplicate Device(s) on Target.....In Progress.
Duplicate Device(s) on Target.....Done.
Create Initiator Group(s) on Target.....Started.
Create Initiator Group(s) on Target.....Done.
Create Port Group(s) on Target.....Started.
Create Port Group(s) on Target.....Done.
Start Data Replication.....Started.
Start Data Replication.....Done.
```

The DM 'Precopy Create' operation successfully executed for storage group 'DM_APP1'.

4. OPTIONAL Ready tgt:

NOTE: Readytgt can only be used after `-precopy` was successfully executed.

This reconfigures the replication so the I/Os are replicated to the source array 643 and will make the target devices visible to the host.

```
symdm readytgt -sid 643 -sg DM_APP1
```

A DM 'ReadyTgt' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
Source SID:000197100643
Target SID:000197100644
Analyze Configuration.....Done.
Preparing Devices for Host discovery.....Started.
Preparing Devices for Host discovery.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.
```

The DM 'ReadyTgt' operation successfully executed for storage group 'DM_APP1'.

5. Use the commit operation to complete the migration process by removing application resources from the source array 643 and releasing resources used for the migration.

```
symdm commit -sid 643 -sg DM_APP1
```

A DM 'Commit' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
Source SID:000197100643
Target SID:000197100644
Analyze Configuration.....Done.
Stop Data Replication.....Started.
Stop Data Replication.....Done.
Remove Masking View(s) on Source.....Started.
Remove Masking View(s) on Source.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....In Progress.
Remove Data Replication.....Done.
Remove Replication Environment.....Started.
Remove Replication Environment.....Done.
```

The DM 'Commit' operation successfully executed for storage group 'DM_APP1'.

Example: Suspending, restarting and recovering the Non-Disruptive Migration session

The following example shows how to suspend, restart or recover a migration. The `symdm sync -stop` command suspends the source-target replication to remove the overhead of synchronizing writes back to the source array. This overhead could interfere when testing

the performance of the application on the new array. Removing this overhead is also beneficial when it is unlikely that a `symdm cancel` action will be issued to migrate the application's storage back to the source array.

The `symdm sync -start` command re-establishes the source-target replication. It can be used to reverse the effect of a `symdm sync -stop` action, or after replication has been halted by a link failure.

The `symdm recover` command can be used after correcting the cause of a failed `symdm` action (`create`, `readytgt`, `commit`, or `cancel`) to put the migration into the appropriate state and then repeat or resume the failed action.

1. The `symdm sync -stopoperation` suspends the replication from the source array 643.

```
symdm sync -sid 643 -sg DM_APP1 -stop
```

```
A DM 'Sync Stop' operation is
in progress for storage group 'DM_APP1'. Please wait...
```

```
Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
Analyze Configuration.....Done.
Stop Data Replication.....Started.
Stop Data Replication.....Done.
```

```
The DM 'Sync Stop' operation successfully executed for
storage group 'DM_APP1'.
```

2. The `symdm sync -startoperation` re-establishes replication from the source array 643.

```
symdm sync -sid 643 -sg DM_APP1 -start
```

```
A DM 'Sync Start' operation is
in progress for storage group 'DM_APP1'. Please wait...
```

```
Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
Analyze Configuration.....Done.
Start Data Replication.....Started.
Start Data Replication.....Done.
```

```
The DM 'Sync Start' operation successfully executed for
storage group 'DM_APP1'.
```

3. The `symdm recoveroperation` fixes the migration state and attempts to repeat or resume a failed action:

- Recovering a failed create action:

```
symdm recover -sid 643 -sg DM_APP1
```

```
A DM 'Recover' operation is
in progress for storage group 'DM_APP1'. Please wait...
```

```
Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
Analyze Configuration.....Done.
Initialize Replication Environment.....Not Needed.
Create Storage Group(s) on Target.....Not Needed.
Duplicate Device(s) on Target.....Not Needed.
Create Initiator Group(s) on Target.....Not Needed.
Create Port Group(s) on Target.....Not Needed.
Start Data Replication.....Started.
Start Data Replication.....Done.
Create Masking View(s) on Target.....Started.
Create Masking View(s) on Target.....Done.
```

```
The DM 'Recover' operation successfully executed for
storage group 'DM_APP1'.
```


- Recovering a failed precopy action:

```
symdm recover -sid 643 -sg DM_APP1
```

A DM 'Recover' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
Analyze Configuration.....Done.
Initialize Replication Environment.....Not Needed.
Create Storage Group(s) on Target.....Not Needed.
Duplicate Device(s) on Target.....Not Needed.
Create Initiator Group(s) on Target.....Not Needed.
Create Port Group(s) on Target.....Not Needed.
Start Data Replication.....Started.
Start Data Replication.....Done.
```

The DM 'Recover' operation successfully executed for storage group 'DM_APP1'.

- Recovering a failed commit action:

```
symdm recover -sid 643 -sg DM_APP1
```

A DM 'Recover' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
Analyze Configuration.....Done.
Stop Data Replication.....Not Needed.
Remove Masking View(s) on Source.....Not Needed.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Replication Environment.....Started.
Remove Replication Environment.....Done.
```

The DM 'Recover' operation successfully executed for storage group 'DM_APP1'.

- Recovering a failed cancel action:

```
symdm recover -sid 643 -sg DM_APP1
```

A DM 'Recover' operation is in progress for storage group 'DM_APP1'. Please wait...

```
Analyze Configuration.....Started.
  Source SID:000197100643
  Target SID:000197100644
Analyze Configuration.....Done.
Stop Data Replication.....Not Needed.
Remove Masking View(s) on Target.....Not Needed.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Port Group(s) on Target.....Started.
Remove Port Group(s) on Target.....Done.
Remove Initiator Group(s) on Target.....Started.
Remove Initiator Group(s) on Target.....Done.
Remove Duplicate Device(s) on Target.....Started.
Wait for deallocation to complete.....Started.
Wait for deallocation to complete.....Done.
```

```

Remove Duplicate Device(s) on Target.....Done.
Remove Storage Group(s) on Target.....Started.
Remove Storage Group(s) on Target.....Done.
Remove Replication Environment.....Started.
Remove Replication Environment.....Done.

```

The DM 'Recover' operation successfully executed for storage group 'DM_APP1'.

Example: Cancelling the Non-Disruptive Migration session

The following example shows how to use the `syndm cancel` to cancel a migration after it is created but prior to cutover, and how to use the `syndm cancel -revert` command to cancel a migration after it is created and cutover, but before it is committed. Refer to [Non-Disruptive Migration control actions and dependent migration states](#) on page 386 for migration states that allow the cancel operation.

- Cancel operation prior to cutover. HR2 migration is currently in the Created state and a decision is made not to complete the migration.

- A cancel operation is run specifying target array 015 and storage group HR2.

During the cancel operation the NDM process:

- Severs the replication pathway connection between the source and target devices.
- Removes the storage provisioned on the target array for HR2 by the create operation.

```
syndm cancel -sid 015 -sg HR2
```

A DM 'Cancel' operation is in progress for storage group 'HR2'. Please wait...

```

Analyze Configuration.....Started.
  Source SID:000194902222
  Target SID:000197300015
Analyze Configuration.....Done.
Remove Masking View(s) on Target.....Started.
Remove Masking View(s) on Target.....Done.
Remove Data Replication.....Started.
Remove Data Replication.....Done.
Remove Port Group(s) on Target.....Started.
Remove Port Group(s) on Target.....Done.
Remove Initiator Group(s) on Target.....Started.
Remove Initiator Group(s) on Target.....Done.
Remove Duplicate Device(s) on Target.....Started.
Wait for deallocation to complete .....Started.
Wait for deallocation to complete .....Done.
Remove Duplicate Device(s) on Target.....Done.
Remove Storage Group(s) on Target.....Started.
Remove Storage Group(s) on Target.....Done.

```

The DM 'Cancel' operation successfully executed for storage group 'HR2'

- After the cancel operation is complete, a host rescan that includes the option to remove the dead paths to the target array. This is followed by a multipath reload command.
- Cancel operation after cutover. HR2 migration has been created, and a host rescan and cutover operation have been run. HR2 is currently in the CutoverSync state and a decision is made not to complete the migration.

- A cancel operation, with the `-revert` flag, is run specifying target array 015 and storage group HR2.

During the cancel and revert operation the NDM process:

- Checks that all HR2 data is available on the source array and moves HR2 processing back to the source array.
- Removes the replication pathway connection between the source and target devices.
- Removes the storage provisioned on the target array for HR2 by the create operation.

```
syndm cancel -sid 015 -sg HR2 -revert
```

A DM 'Cancel Revert' operation is in progress for storage group 'DM_APP1'. Please wait...

```

Analyze Configuration.....Started.
  Source SID:000194902222

```

```

Target SID:000197300015
Analyze Configuration.....Done.
Revert Data Replication.....Started.
Revert Data Replication.....In Progress.
Revert Data Replication.....Done.
Remove Masking View(s) on Target.....Started.
.....

```

The DM 'Cancel' operation successfully executed for storage group 'HR2'

2. After the cancel operation is complete, a host rescan that includes the option to remove the dead paths to the target array. This is followed by a multipath reload command.

Application environment following the migration cancellation.

- HR2 continues to run uninterrupted on the source array as the migration process cancels the migration.
- At the completion of the cancel operation, the HR2 environment is the same as it was prior to running the migration create operation. No resources remain allocated for HR2 on the target array, and HR2 is processing only on the source array.

Example: Troubleshoot migration failure (device moved between storage groups)

The following example shows how to use the `symdm list -v` command to help troubleshoot a failed migration session. The failure is a result of a device being moved, in the target configuration, from one child storage group to another after the migration was started.

The following output from the `symdm list -v` command displays the migration status of the application HR2 from array 137 to array 084. The application contains a cascaded storage group (two child storage groups) with three masking views to the application hosts. These masking views are configured with different port groups and a combination of cascaded and standalone initiator groups.

1. Run the `symdm list -v` command to display summary information for the storage group HR2 migration session:

```

symdm list -sid 084 -sg HR2 -v

Storage Group      : HR2
Source Array       : 000194901137
Target Array       : 000197100084

Migration State    : CutoverReady
Total Capacity (GB) : 500.0
Done (%)           : N/A

Source Configuration: OK
{
  Storage Groups (3) : OK
  Masking Views (3)  : OK
  Initiator Groups (5): OK
  Port Groups (3)    : OK
}

Target Configuration: Failed
{
  Storage Groups (3) : Failed
  Masking Views (3)  : OK
  Initiator Groups (5): OK
  Port Group (3)     : OK
}

Device Pairs (9): OK

```

The output shows a migration failure for the storage group objects in the target configuration. The migration has reached *CutoverReady* state.

2. Run the `symdm list -sg -v -detail` with the `-sg_info` option. This filters the display to list only the storage groups used in the migration session with storage group HR2, and the current state of each groups.

```

symdm list -sid 084 -sg HR2 -v -detail

Storage Group      : HR2
Source Array       : 000194901137
Target Array       : 000197100084

```

```
Migration State      : CutoverReady
Total Capacity (GB) : 500.0
Done (%)            : N/A
```

```
Source Configuration: OK
```

```
{
  Storage Groups (3): OK
```

```
{
  Name      : HR2
  Status    : OK
  {
    Group Name
    -----
    HR2_sga
    HR2_sgb
  }
```

```
Name      : HR2_sga
Status    : OK
{
  Dev
  -----
  002B7:002B9
}
```

```
Name      : HR2_sgb
Status    : OK
{
  Dev
  -----
  0030A 003B2:003B6
}
```

```
Target Configuration: Failed
```

```
{
  Storage Groups (3): Failed
```

```
{
  Name      : HR2
  Status    : OK
  {
    Group Name
    -----
    HR2_sga
    HR2_sgb
  }
```

```
Name      : HR2_sga
Status    : Failed
{
  Dev
  -----
  01F28:01F2A
}
```

```
Name      : HR2_sgb
Status    : Failed
{
  Dev
  -----
  01F2B:01F30
}
```

The child storage groups in the target configuration are in a *Failed* state because they no longer contain the same devices as the storage groups in the target array (device was moved).

3. Identify the moved device by running the `symmsg show` command on the target array, then compare these devices with the expected set of devices shown in the display. Once the device discrepancy is corrected, a `symdm cutover` command can be run to repeat the failed action.

Example: Troubleshoot migration failure (masking view added)

The following example shows how to use the `symdm list -v` command to help troubleshoot a failed migration session. The failure is a result of a masking view being added, in the source configuration, after the migration was started.

The following output from the `symdm list -v` command displays the migration status of the application HR2 from array 137 to array 084. The application contains a cascaded storage group (two child storage groups) with three masking views to the application hosts. These masking views are configured with different port groups and a combination of cascaded and standalone initiator groups.

1. Run the `symdm list -v` command to display summary information for the storage group HR2 migration session:

```
symdm list -sid 084 -sg HR2 -v
Storage Group      : HR2
Source Array       : 000194901137
Target Array       : 000197100084

Migration State    : CutoverSync
Total Capacity (GB) : 500.0
Done (%)           : N/A

Source Configuration: Failed
{
  Storage Groups (3) : OK
  Masking Views (3)  : Failed
  Initiator Groups (5): OK
  Port Groups (3)    : OK
}

Target Configuration: Failed
{
  Storage Groups (3)      : OK
  Masking Views (3)       : OK
  Initiator Groups (5)    : OK
  Port Group (3)          : OK
}

Device Pairs (9): OK
```

The output shows a migration failure for the masking view objects in the source configuration. The migration has reached *CutoverSync* state.

2. Run the `symdm list -sg -v -detail` command to display a detailed list of all the objects used in the migration session with storage group HR2, and the current state of each object.

 **NOTE: Output is abbreviated. Also, the `-view_info` option can be used to display only the masking view info.**

```
symdm list -sid 084 -sg HR2 -v -detail

Storage Group      : HR2
Source Array       : 000194901137
Target Array       : 000197100084

Migration State    : CutoverSync
Total Capacity (GB) : 500.0
Done (%)           : 100

Source Configuration: Failed
{
  Storage Groups (3): OK
  {
    Name      : HR2
    Status    : OK
    {
      Group Name
      -----
      HR2_sga
      HR2_sgb
      ....
    }
  }
}
```

```

Masking Views (3): Failed
{
  Masking View Name      Initiator Group      Port Group      Status
  -----
  HR2_hosta_view         HR2_hosts            HR_pg1          OK
  HR2_hostb_view         HR2_hostx_ig         HR_pg2          OK
  HR2_hostc_view         HR2_hosty_ig         HR_pg3          OK
}

Initiator Groups (5): OK
{
  Name      : HR2_hosts
  Status    : OK
  {
    Group Name
    -----
    HR2_hostx_child_ig
    HR2_hosty_child_ig
  }
  ....
}

Port Groups (3): OK
{
  Name      : HR2_pg1
  Status    : OK
  {
    Dirport Status
    -----
    02E:000 OK
    03F:000 OK
    02F:000 OK
    03G:000 OK
  }
  ....
}

Target Configuration: OK
{
  Storage Groups (3): OK
  {
    Name      : HR2
    Status    : OK
    {
      Group Name
      -----
      HR2_sga
      HR2_sgb
    }
    ....
  }

  Masking Views (3): OK
  {
    Masking View Name      Initiator Group      Port Group      Status
    -----
    HR2_hosta_view         HR2_hosts            HR_pg1          OK
    HR2_hostb_view         HR2_hostx_ig         HR_pg2          OK
    HR2_hostc_view         HR2_hosty_ig         HR_pg3          OK
  }

  Initiator Groups (5): OK
  {
    Name      : HR2_hosts
    Status    : OK
    {
      Group Name
      -----
      HR2_hostx_child_ig
      HR2_hosty_child_ig
    }
    ....
  }

  Port Group (3): OK
  {
    Name      : HR2_pg1
  }
}

```

```

Status : OK
{
  Dirport Status
  -----
  07E:010 OK
  06E:011 OK
  ....
Device Pairs (9): OK
{
  Source      Target
  Dev        Status Dev        Status
  -----
  002B7 OK      01F28 OK
  002B8 OK      01F29 OK
  002B9 OK      01F2A OK
  0030A OK      01F2B OK
  003B2 OK      01F2C OK
  003B3 OK      01F2D OK
  003B4 OK      01F2E OK
  003B5 OK      01F2F OK
  003B6 OK      01F30 OK

```

The masking views in the source configuration are in a *Failed* state because the source array no longer contains the same list of masking views as the target array (a masking view was added).

3. Identify the added masking view by running the `symaccess show HR2 -type storage` command, then compare the current list of masking views on the storage group with the list of masking views shown in the display. Once the masking view discrepancy is corrected, a `symdm commit` command can be run to repeat the failed action.

Open Replicator

This chapter summarizes the Open Replicator feature which can be used to migrate data from older arrays and some third party arrays to arrays running HYPERMAX OS 5977 and higher.

The Open Replicator feature and **symrcopy** command used to migrate data is documented in the Dell EMC Solutions Enabler Array Controls and Management CLI User Guide version 8.3 and higher.

Topics:

- [Open Replicator and arrays running HYPERMAX OS](#)
- [ORS and host interaction](#)
- [ORS rcopy concepts](#)
- [ORS operational rules and limitations](#)
- [ORS copying limitations](#)
- [ORS device guidelines](#)
- [ORS SAN setup requirements](#)
- [ORS SYMCLI symsan support](#)
- [Open Replicator session options](#)
- [Open Replicator control operations](#)
- [Open Replicator examples](#)
- [Open Replicator operational restrictions and state reference](#)

Open Replicator and arrays running HYPERMAX OS

The Open Replicator (ORS) **symrcopy** command provides a method for copying data to or from various types of arrays within a storage area network (SAN) infrastructure. For example, Open Replicator provides a tool used to migrate data from older arrays, and some third-party storage arrays to VMAX arrays including arrays running HYPERMAX OS 5977.

The ORS feature and **symrcopy** command, used to migrate data, is documented in the Dell EMC Solutions Enabler Array Controls and Management CLI User Guide version 8.3 and higher.

The following ORS rules and limitations apply to control arrays running HYPERMAX OS 5977 or higher:

- ORS push sessions are not supported, along with differential copying, precopy option, recreate, restore, and remove command.
- The maximum number of active sessions allowed is 512 when running HYPERMAX OS 5977 or higher.
- Only one remote array is supported.
- Creating ORS sessions on control devices that are involved in a TimeFinder session or a SRDF operation is not supported.
- Any VMAX3 array FA port, connected to the remote array and granted access to the remote devices, can be used for Open Replicator sessions.
- Requires at least one zoned remote port for cold sessions, and it is recommended that there are at least two zoned remote ports for hot sessions, although this recommendation is not enforced.
- Set ceiling can be set to DISABLE.
- Dynamically finds directors and ports that can access an ORS session's remote devices, and uses these directors and ports for data transfer.
- Setting session pace is not supported.

ORS and host interaction

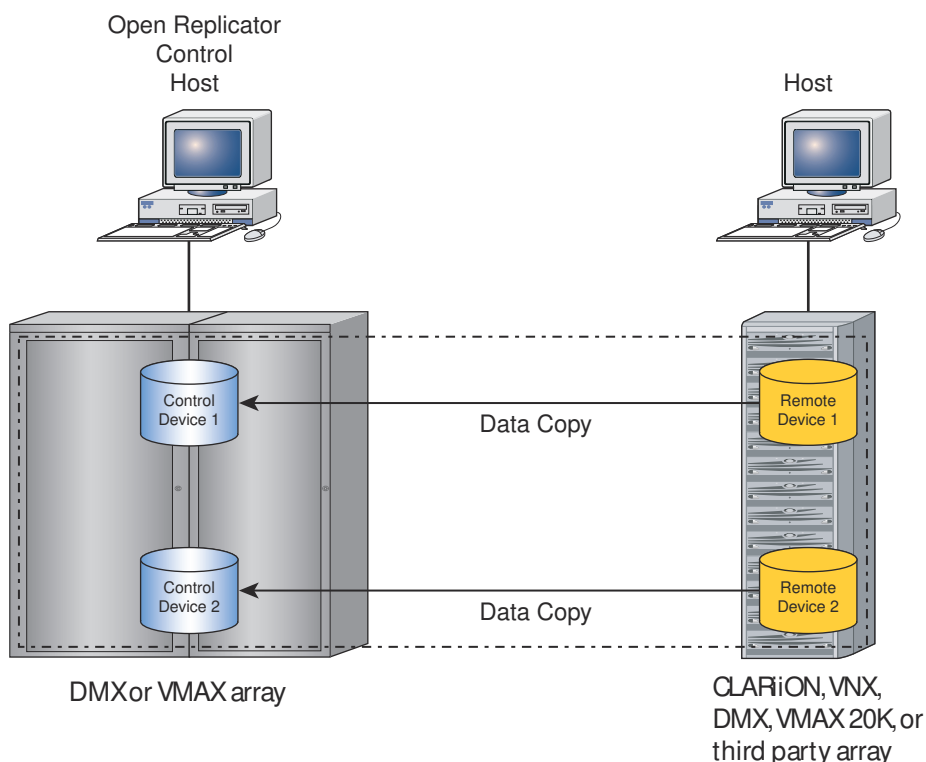
Open Replicator copy (Rcopy) operations are controlled from a local host attached to the DMX or VMAX Family array. Data copying is accomplished as part of the storage system process and does not require host resources. The data can be copied online between DMX and VMAX arrays allowing host applications, such as a database or file server, to remain operational (function normally) during the copy process.

ORS rcopy concepts

The following Rcopy concepts and terminology are used throughout the Open Replicator Migration section of this product guide:

- The VMAX and DMX arrays and its devices are referred to as the *control* side of the copy operation. Older DMX or VMAX arrays, or third-party arrays on the SAN are referred to as the *remote* array/devices.
- The copy direction is from the perspective of the control side. There are two types of copy operations, *push* and *pull*. A push operation copies data from the control device to the remote device(s). A pull operation copies data to the control device from the remote device(s).
- Copy operations are either *hot* (online) or *cold* (offline).
- Use the `-name` option to give the session a name. Use the `-session_name` option when specifying the session name for control operations.
- There can be only one control device per active session.

Open Replicator can be used to migrate data into a VMAX Family array from older DMX and VMAX arrays, or other third-party arrays. [Figure 15. Control array device pull operation](#) on page 417 shows two Open Replicator copy sessions performing a pull operation, where data is copied through the SAN infrastructure from remote devices to the control array.



SYM-001467

Figure 15. Control array device pull operation

NOTE:

Since data is copied through the SAN infrastructure, Open Replicator may require updating the zoning configuration before copying data between arrays is allowed. For zoning requirements and suggestions, refer to [ORS SAN setup requirements](#) on page 419.

Open Replicator can be used to copy data from a control array to older array. [ORS rcopy concepts](#) on page 417 shows two Open Replicator copy sessions performing a push operation, where data is copied from the control array to remote devices within the SAN infrastructure.

NOTE:

For arrays running HYPERMAX OS 5977, push operations are not supported.

ORS operational rules and limitations

The following general rules apply to Open Replicator sessions:

- Remote devices do not have to be the same RAID type or meta-configuration.
- On pull operations, the remote devices should not be updated by array hosts for the duration of the copy process.
- For pull operations from devices with SCSI reservations, if the remote devices have a cluster running against them or the devices are AIX LVM devices, the cluster, AIX host, or other software that is creating the SCSI reservations must be shutdown before creating the Open Replicator session.
- Data corruption to devices is possible during a copy operation if another host on the SAN has write access to the remote device. EMC recommends that the remote device be unmounted or marked as Not Ready to any other hosts on the SAN to guarantee that the device cannot change while copying is in process.
- Accumulated I/O errors between the control device and remote device will cause a session to fail if the copy operation is a hot push. The failed session may be activated again as long as no new data has been written to the control device since the session failed. The session will temporarily stall and restart on any other type of copy operation.
- Open Replicator fully supports copy operations for thin devices. For information on Virtual Provisioning and creating thin devices, refer to [Thin Device Management](#) on page 71.

ORS copying limitations

Copying is device-based; extent copying is not supported. Device configuration changes cannot be made during an Open Replicator session, as making device changes may lead to inconsistent data on the local device if pulling, or on the remote device if pushing data.

Open Replicator cannot detect changes to a remote device during, or between incremental copies. Before each session, make sure that there are no changes being made to the remote device.

ORS device guidelines

Table 26. Control and remote device guidelines

Action	Control device	Remote device
Creating the device file	DMX or VMAX Family arrays Always listed on left Format: symdev=arrayid:device Example: symdev=7098:E9	DMX or older VMAX array, or third-party array Always listed on right Format: symdev =array:device or wwn=WWN Example: wwn=6006048000000000314353594D303737 Refer to ORS and creating a device file on page 422 for device format rules
Hot push (not supported on HYPERMAX OS 5977)	One device per session All directors must see remote device	One device per session Device not accessible to host Can use -consistent , -nodifferential , and -precopy NOTE: Differential copying is the default for copy sessions and does not need to be specified in the <code>symrcopy create</code> command.
Cold push (not supported on HYPERMAX OS 5977)	One device per session Device Not Ready to host	Can use -nodifferential NOTE:

Table 26. Control and remote device guidelines (continued)

Action	Control device	Remote device
	At least one director must see remote device(s)	Differential copying is the default for copy sessions and does not need to be specified in the <code>symrcopy create</code> command.
Hot pull	One device per session Device online to the host All directors must see the remote device ^a	One device per active session Can use <code>-donor_update,-frontend_zero</code>
Cold pull	One device per session At least one director must see the remote device and the devices is not ready to host	One device per session Can use <code>-frontend_zero</code>

a. Applies only to devices running on arrays Engenuity 5876 or lower. This is not a requirement for HYPERMAX OS 5977 or higher.

ORS SAN setup requirements

Since data is copied through the SAN infrastructure, Open Replicator may require a SAN configuration update before copying data between storage arrays is allowed. Because of the various types of cabling, zoning, and masking that can exist within a SAN configuration, the following requirements are provided as a generic reference for setting up a data migration with Open Replicator through a SAN:

- A Fibre Channel switch is required for Open Replicator. Direct connections (such as arbitrated loop) are not supported.
- For arrays running Engenuity 5773 and 5876:
 - The SAN for the remote storage array must have connectivity to the control array SAN. Open Replicator requires that at least one port on the remote array, that allows access to the remote device, have access to the control device through at least one port for a cold copy and all ports for a hot copy on the control array.
 - Zoning must be set up on the Fibre Channel switch to zone the control array fibre adapters (FAs) to the remote storage array front-end adapter(s).
 - If the storage port for the source devices is running volume configuration management (VCM) software such as Volume Logix for Symmetrix or a similar software product, setup may require granting permission to access the target device(s) from the target storage port. For instance, on a VMAX array, the control FA(s) must be enabled to have access to the remote storage array device(s).
 - Specific instructions for granting access to target devices vary depending on the selected target storage array and VCM software.
- For arrays running HYPERMAX OS 5977:
 - Any VMAX3 array FA port, connected to the remote array and granted access to the remote devices, can be used for Open Replicator sessions. This is a change from earlier versions of ORS where the ports involved in Open Replicator sessions had to be FA ports where the volumes were mapped.
 - Requires at least one zoned remote port for cold sessions, and it is recommended that there are at least two zoned remote ports for hot sessions. The recommendation for hot sessions is not enforced.

ORS SYMCLI symsan support

The SYMCLI command **symsan** lists port and LUN WWNs as seen from a specific array director and port. This is used to validate that the zoning between the port and target is correct. It does not require a created Open Replicator session. Use this command to display remote port WWNs, and LUN WWNs seen behind a remote port WWN.

The **symsan** command allows for:

- Listing all ports or LUNS in the SAN that are seen by a specific DX director or all DX directors.
- Listing all ports or LUNS in the SAN that are seen by a specific FA director or all FA directors.
- Listing all ports or LUNS in the SAN that are seen by a specific FA/DX director or all FA/DX directors, using the `-dir` option.

Refer to [SYMCLI help options](#) on page 24 for information on how to access `symsan` manpage or command help.

Open Replicator session options

Open Replicator copies data in sessions across the SAN infrastructure. A device file is used to specify the device pairs to be used in the copy session. These devices are referred to as the control and remote devices. The control device always resides on the locally-attached DMX or VMAX Family array, and is responsible for controlling data copying to or from its partner remote device. Devices listed in the device file are identified by either logical unit number (LUN), World Wide Name (WWN), or by a combination of the storage array ID and device name (use `symdev` for arrays). Refer to [ORS and creating a device file](#) on page 422 for instructions on how to obtain device information and create the device file.

A copy session is first defined by using the `symrcopy create` command. A session name can be specified for later use in control operations. The push/pull options (`-push` | `-pull`) define the direction of the copy operation for device pairs listed in the device file, and the hot/cold options (`-hot` | `-cold`) define online or offline copying.

symrcopy session options

`-pull`

Pulls data through the SAN to the control device(s) from the remote device(s).

`-push`

Pushes data across the SAN from the control device(s) to the remote device(s)

`-hot`

Control device is available as read/write online to the host while the copy operation is in progress. With hot copying, 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.

`-cold`

Control device is unavailable 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.

NOTE:

Arrays running HYPERMAX OS 5977 dynamically find directors and ports that can access an ORS session's remote devices, and uses these directors and ports for data transfer.

If a control device is pushing data to a remote device and that control device is currently online for host write I/O operations, a consistent point-in-time copy can be made across multiple control devices using the Enginuity Consistency Assist (ECA) feature (`-consistent`). This will temporarily prevent any host write I/Os while the Open Replicator copy session begins.

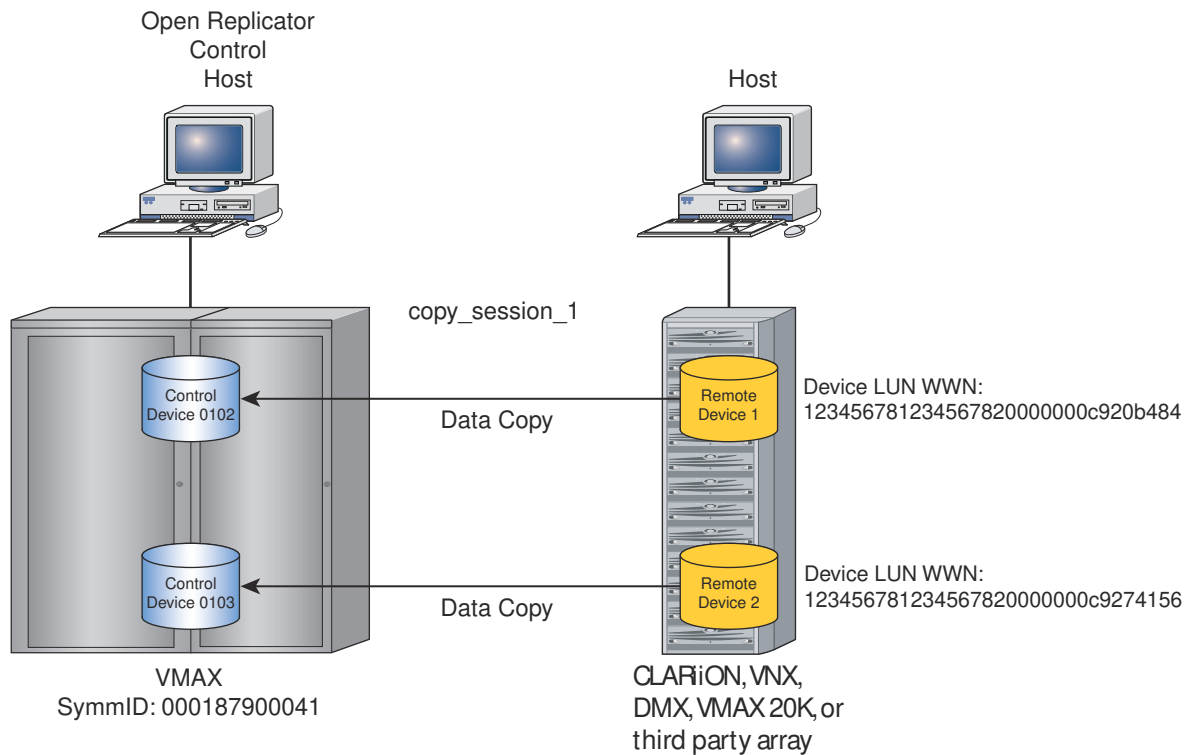
ORS hot pull copy session example

Figure 16. Control array hot pull using the `symrcopy` command on page 421 shows Open Replicator copy sessions created and activated for a hot pull copy operation. A device file `-file filename` contains the pairing information for the control and remote devices. Each line in the device file is a copy session. The file specifies control devices by "Symmetrix ID: device number" and remote devices by "LUN WWN" as follows:

```
symdev=000187900041:0102   wwn=123456781234567820000000c920b484
symdev=000187900041:0103   wwn=123456781234567820000000c9274156
```

NOTE:

Refer to [ORS and creating a device file](#) on page 422 for instructions on how to obtain device information and create the device file.



```
symrcopy create -name copy_session_1 -pull -hot -file pairs
symrcopy activate -file pairs
```

Figure 16. Control array hot pull using the symrcopy command

Open Replicator control operations

SYMCLI Open Replicator performs control operations from a local host attached to the DMX or VMAX Family array and implements these operations in sessions across the SAN infrastructure. Open Replicator copy sessions are first created using a device file, which lists the device pairs (control and remote) for the operation.

NOTE: For arrays running HYPERMAX OS 5977, only one remote target is supported.

Open Replicator SYMCLI **symrcopy** command performs the copy sessions. The main control operations, required to successfully complete a copy session, are as follows:

- Create the session.
- Activate the session.
- Terminate the session.

HYPERMAX OS 5977 settings and actions

- Data protection and recovery options for hot and cold pulls.
- Enable or disable front-end zero detection for pull operations to thin control devices.
- Background copying mode of a session.
- Ceiling value for bandwidth.

NOTE:

Ceiling value is set using the **symqos -rcopy** command.

- List, query, and verify copy sessions to display the current session status.
- Export the run information to an output file.

Engenuity 5876 or lower settings and actions

- Data protection and recovery options for hot pulls.
- Enable or disable front-end zero detection for pull operations to thin control devices.
- Background copying mode of a session.
- Ceiling value for bandwidth.

NOTE:

Ceiling value is set using the `symqos -rcopy` command.

- Session pace for copying, recreating, and restoring.
- Differential copying.
- List, query, and verify copy sessions to display the current session status.
- Remove a remote device from a session.
- Recreate a differential copy session.
- Rename a differential copy session.
- Restore data from a remote device of a copy session.
- Export the run information to an output file.

NOTE:

For detailed syntax of the `symrcopy` command, refer to the *EMC Solutions Enabler CLI Command Reference*

ORS and creating a device file

Before an Open Replicator copy session can be created, a device file must be created that lists the control and remote device pairs for the copy operation. The device file syntax contains two columns (one for control devices and one for remote devices). Devices in the file are specified either by their unique LUN WWN, or by the storage array ID and device number (Storage ID:device#).

Use the following rules to determine the correct device ID format in the device file:

- If the array for the remote device is visible to the host where the `symrcopy` command is run (locally or by a remote RDF connection), either the storage array ID and device number or the LUN WWN can be used for the remote device ID in the device file.
- If the array for the remote device is only visible using the `symsan` command, only the LUN WWN can be used for the remote device ID in the device file.

Obtain device information

Description

The Solutions Enabler (SYMCLI) uses several commands to obtain device information, including device number, director information, WWN, and capacity. This information is helpful in determining and identifying devices for use in a device file. Some of these commands include: `symdev`, `syminq`, `sympd`, `symsan`.

Examples

To list of devices on array 041, enter:

```
symdev list -sid 0041
```

To list devices by device wwn (a SCSI inquiry), enter:

```
syminq -sym -wwn
```

Sample output

For array 041 devices:

```
Symmetrix ID: 000187900041
  Device Name          Directors          Device
-----
Sym  Physical          SA :P DA :IT  Config          Attribute  Sts  Cap
-----
                                     (MB)
```

```

00102 /dev/vx/rdmp/c5t0d2s2 03A:0 01A:C2 2-Way Mir Grp'd (M) RW 17261
00103 /dev/vx/rdmp/c5t0d3s2 03A:0 01D:C3 2-Way Mir Grp'd (M) RW 17261
00104 /dev/vx/rdmp/c5t0d4s2 03A:0 16C:D2 TDEV N/Grp'd (M) RW 17261
00105 /dev/vx/rdmp/c5t0d5s2 03A:0 16C:C3 TDEV N/Grp'd (M) RW 17261
00106 /dev/vx/rdmp/c5t0d6s2 03A:0 01A:C4 RAID-5 Grp'd (M) RW 17261
00107 /dev/vx/rdmp/c5t0d7s2 03A:0 01A:D5 RAID-5 Grp'd (M) RW 17261
< . . . >

```

For device wwn list:

```

syminq -sym -wwn
Device
-----
Name Num Array ID WWN
-----
/dev/vx/rdmp/c5t0d2s2 00168 000000006190 60060480000190300016533030303238
/dev/vx/rdmp/c5t0d3s2 001F8 000000006190 60060480000190300016533030303239
/dev/vx/rdmp/c5t0d4s2 001F9 000000006190 60060480000190300016533030303241
/dev/vx/rdmp/c5t0d5s2 00170 000000006190 60060480000190300016533030303242
/dev/vx/rdmp/c5t0d6s2 00172 000000006190 60060480000190300016533030303243
/dev/vx/rdmp/c5t0d7s2 001B2 000000006190 60060480000190300016533030303244

```

NOTE:

For details about the `symdev` and `syminq` commands syntax, refer to the *Dell EMC Solutions Enabler CLI Command Reference*

Create ORS device file

The control device is always listed in the first column of the device file. Lines in the device file that begin with a pound symbol (#) will be ignored. The device filename (`-file Filename`) is inserted into the command line for control operations.

This example shows a device file with multiple copy sessions. Each line in the device file is a separate copy session.

```

# dev_file_1
## column1:control column2:remote
# Symmetrix and StorageID:device always listed first
symdev=000000006190:0168 wwn=6006048000000000619053594D314638
symdev=000000006190:01F8 wwn=6006048000000000619053594D314640
symdev=000000006190:01F9 wwn=6006048000000000619053594D314637
symdev=000000006190:0170 wwn=6006048000000000619053594D314642
symdev=000000006190:0172 wwn=6006048000000000619053594D314646
symdev=000000006190:01B2 wwn=6006048000000000619053594D314649
# End

```

This example shows a device file with one control device (01) and multiple remote devices (41 and 42). This is only used with cold push sessions:

```

symdev=000000001234:01 symdev=000000005678:41
symdev=000000001234:01 symdev=000000005678:42

```

This example shows a device file with a mix of `symdev` and `wwn` device IDs.

```

symdev=000000001234:01 symdev=000000005678:42
symdev=000000001234:02 symdev=000000005678:43
symdev=000000001234:03 wwn=6006048000000000567853594D303434

```

Export ORS device list to text file

Examples

To export session device list to text file `dev_file_1.txt`, enter:

```

symrcopy export -session_name rcopy_1 -file dev_file_1.txt

```

Sample output

The output file (`dev_file_1.txt`) contains the session device list.

```
symdev=000000006190:0168 wwn=600604800000000619053594D314638
symdev=000000006190:01F8 wwn=600604800000000619053594D314640
symdev=000000006190:01F9 wwn=600604800000000619053594D314637
symdev=000000006190:0170 wwn=600604800000000619053594D314642
symdev=000000006190:0172 wwn=600604800000000619053594D314646
symdev=000000006190:01B2 wwn=600604800000000619053594D314649
```

Create ORS copy session

Description

The `symrcopy create` defines a new ORS copy session. Other mandatory syntax session controls included in the `symrcopy create` command line are the copy direction parameter (`-pull`), the online/offline parameter (`-hot`), and the device text filename (`-file FileName`). See [Open Replicator session options](#) on page 420 for more information on the mandatory parameters.

Syntax

```
symrcopy create -name <SessionName> <-pull> <-hot> -file <FileName>
```

Options

`-name`

Session name — used for control operations.

Examples

To define a hot, pull Open Replicator copy session named `rcopy_1`, enter:

```
symrcopy create -name rcopy_1 -pull -hot -file dev_file_1
```

NOTE: Arrays running HYPERMAX OS 5977 dynamically find directors and ports that can access an ORS session's remote devices, and uses these directors and ports for data transfer.

ORS front-end zero detection

Front-end zero detection is an option used with thin control devices. Front-end zero detection looks for incoming zero patterns from the remote device, and instead of writing the incoming data of all zeros to the thin control device, the group on the thin device is de-allocated. Front-end zero detection is allowed for pull operations only and is indicated by the `frontend_zero` option. This option is ignored for any standard (thick) control devices.

NOTE:

With front-end zero detection enabled, "persistent" allocations are treated as regular allocations, and track group is de-allocated.

Enable front-end zero detection

Example

By default pull sessions are disabled for the `-frontend_zero` detection option. To create a hot pull session and enable front-end zero detection, enter:

```
symrcopy create -session_name rcopy_1 -pull -hot -frontend_zero -file dev_file_1
```


Disable front-end zero detection

Examples

The `set frontend_zero off` option disables zero detection for the session, and is only allowed during active `-pull` operations. To disable front-end zero detection for the active session `rcopy_1`, enter:

```
symrcopy -session_name rcopy_1 set frontend_zero off
```

NOTE: Once front-end zero detection is disabled during an active session, it cannot be enabled again during the session.

ORS donor update

To protect against potential data loss due to a SAN failure or other connectivity issue during a hot pull operation, use the `-donor_update` option. With this option, all writes to the control device from the host are immediately copied to the remote device as well. Because the data is fully copied to both the remote device and the control device, if a failure occurs, the session can be safely terminated and created again to fully recover from any mid-copy failure.

Enable ORS donor update

Examples

To create and activate an Open Replicator copy session for a hot pull operation using the `donor_update` option:

```
symrcopy create -name rcopy_1 -pull -hot -donor_update -file dev_file_1
symrcopy activate -session_name rcopy_1
```

NOTE: For information on the `activate` command, refer to [Activate ORS session on page 426](#).

Terminate and restart ORS hot pull session

Examples

If during an activated hot pull operation, a SAN failure or other connectivity issue is detected, then terminate the Open Replicator sessions. To terminate an Open Replicator sessions associated with the control device:

```
symrcopy terminate -file dev_file_1 -symforce
```

To start the copy session again after the problem is resolved, and restart the copy process from the point of failure:

NOTE: The `donor_update` option must be included in the original `symrcopy create` command to fully recover all writes made to the devices prior to the failure.

```
symrcopy create -name rcopy_1 -pull -hot -donor_update -file dev_file_1
symrcopy activate -session_name rcopy_1
```

NOTE: The above example restarts the copy process from where it left off at the time of failure.

Disable ORS donor update

Description

The `set donor_update off` option disables the `donor_update` option. This command stops the copying of data to the remote devices, and stops all new writes to the control device from immediately being copying to remote device. The donor update option may also be used on an incremental restore session. Refer to [Restoring an ORS session on page 437](#) for more information.

NOTE: The `donor_update` option may also be turned off while the session is in the *CopyInProg* (copy in progress) state by including the `-force` option in the command line. The session will continue to copy in its current mode without donor update.

Options

-consistent

Maintains consistency of the data on the remote device. Without this option, donor update is still deactivated, but consistency on the remote devices is not maintained. This option is useful for a hot pull session, with donor update enabled, where the session has finished copying and maintaining a consistent image on the remote devices is desired. By using the `set donor_update off` command with the `-consistent` option after the session has fully copied, the donor update portion of the session is disabled, but data consistency on the remote devices is maintained.

Examples

To set the donor update option to off and maintain consistency on the remote devices for session `rcopy1`, enter:

```
symrcopy set donor_update off -session_name rcopy_1 -consistent
```

Restrictions:

If the session is terminated, renamed, restored, recreated, a device is removed, or another session is created using the same control device, the donor update portion of the session is automatically deactivated and consistency on the remote devices is lost. To maintain the consistency on the remote devices, issue the `set donor_update off -consistent` command prior to any of these actions.

Activate ORS session

Description

The `symrcopy activate` command activates the copy sessions for device pairs listed in the device file and begins copying data to (pushing) or from (pulling) the remote devices.

If control devices are pushing data to remote devices and the control devices are currently online for host I/O operations, include the Engineuity Consistency Assist (ECA) option (`-consistent`) in the command line to temporarily prevent host I/O while the Open Replicator copy session begins. This begins a consistent point-in-time copy to the remote devices using an ECA window, which temporarily freezes host I/O to the control devices.

Syntax

To begin the copying process for an Open Replicator copy session, use the following syntax:

```
symrcopy activate -file Filename -session_name SessionName
```

NOTE: Any other Open Replicator copy sessions that were previously created using the specified device file (and session name) will also be started.

Options

-file *Filename*

The device file.

-session_name *SessionName*

The session name.

Example

To activate an Open Replicator copy session, enter:

```
symrcopy activate -session_name rcopy_1
```

NOTE: Under certain circumstances, failed sessions may be reactivated. Refer to [Recover from failed ORS session on page 436](#).

ORS background copying

Open Replicator copy sessions that are actively background copying to devices are in the *CopyInProg* state. This is the default state for copy sessions. This state can be changed to the *CopyOnAccess* state for a pull operation, the *CopyOnWrite* state for a push operation, or the *Precopy* state for a hot push operation by using the `symrcopy set mode [-copy|-nocopy|-precopy]` options. An activated session in the *CopyOnAccess* state copies data to the control device only when those tracks have been accessed on the control device. An activated session in the *CopyOnWrite* state copies data to the remote device only when those tracks are accessed on the control device.

A hot push session that is in the *Precopy* state immediately begins copying data in the background before the session is activated. Session data will continue copying to the remote device until either the mode is changed to `nocopy`, `copy`, or the session is activated, at which time a point-in-time copy of the control device is made. After the session has been activated, copying will continue in the *CopyOnWrite* (`nocopy`) or *CopyInProg* (`copy`) state. The *Precopy* feature is available only for hot push operations. Hot push sessions can also be set to *Precopy* mode by including the `-precopy` option with either the `create` or `recreate` command.

The background copy status for a session is designated by a flag in the output of `symrcopy list` command. Refer to [Monitor ORS session status](#) on page 428 for more detail.

NOTE: The `-precopy` option requires Engenuity version 5773 or 5876.

Stop and restart ORS background copying

Examples

To temporarily stop the background copying for a session by changing the state to *CopyOnAccess* or *CopyOnWrite* from *CopyInProg* using the `symrcopy` command, enter:

```
symrcopy set mode nocopy -file dev_file_1
```

To resume background copying for a session and change to the *CopyInProg* state, enter:

```
symrcopy set mode copy -file dev_file_1
```

NOTE: The `-precopy` option requires Engenuity version 5773 or 5876.

ORS background copying (hot) without point-in-time copy

Examples

The following are examples of how to immediately begin background copying on a hot push session without making a point-in-time copy:

```
symrcopy set mode precopy -file dev_file_1
```

```
symrcopy create dev_file_1 -precopy
```

NOTE: To see the `-precopy` option used with the `recreate` command, refer to [Recreate an ORS session](#) on page 435.

Restrictions:

- The `-precopy` option requires Engenuity version 5773 or 5876.
- Precopy mode can only be set when the session is not activated.

ORS ceiling value

Ceiling values can be adjusted to optimize the performance of the specific SAN environment.

The `symqos -rcopy set ceiling` command sets the maximum allowed bandwidth percentage for a given director, port, director/port pair, or all directors and ports. Valid values are 0 - 100 (%), `DISABLE` or `NONE` (shuts off the ceiling function).

For example, setting the ceiling value to 100% causes Open Replicator to consume as much bandwidth as possible, typically:

- 80 MB/s for a 1 GB SAN

- 150 MB/s for a 2 GB SAN
- 180 MB/s for a 4 GB SAN (for DMX - 4 systems)
- 300 MB/s for a 4 GB SAN (for VMAX Family arrays)
- 300 MB/s for an 8 GB SAN (for VMAX Family arrays)

Setting the ceiling to a value (other than NONE) renders the session pace value ineffective to the copy. If the ceiling value is set to NONE, the session pace is in effect for the copy.

When using ORS with SRDF/A, users should monitor and adjust the ORS ceiling/session pace or turn on SRDF/A pacing to prevent ORS I/O from causing SRDF/A to drop. Refer to Primus case emc292509 for guidelines on setting these values.

 **NOTE: The DISABLE setting is only allowed for HYPERMAX OS 5977 or higher.**

Set ORS ceiling

Examples

To set a bandwidth ceiling of 100% for all directors on array 6190, enter:

```
symqos -rcopy set ceiling 100 -dir all -sid 6190
```

To view the ceiling setting, enter:

```
symqos -rcopy list ceiling -dir all -sid 6190
```

Set the ORS session pace

Description

If the ceiling value is set to NONE, the session pace can be set for devices being copied, recreated, or restored to manage the speed of the replication process. The session pace designates how fast data copies between devices. 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.

Example

The following example shows how to set the session copy pace:

```
symrcopy set pace 0 -file dev_file_1
```

Restrictions:

- The session pace is ineffective to the copy if the ceiling is set to a value other than NONE.
- For arrays running HYPERMAX OS 5977 or higher, setting session pace is not supported.

Monitor ORS session status

Open Replicator session status is checked using the `symrcopy query`, `symrcopy list`, or `symrcopy verify` command.

List all ORS sessions

Description

To list the Open Replicator copy sessions for a local array, use the `symrcopy list` command. This command returns status information for all created sessions. To list session information for a specific Symmetrix array, include the array ID (`-sid SymmID`) option in the command line.

Options

The following options are available for the `symrcopy list` command:

-offline

It only displays information held in the database and does not query the array for updated session information.

-detail

It displays additional device information for modified tracks, session pace, and session name.

-wwn

It displays the full device world wide name.

NOTE: Using the -detail and -wwn options expands the width of the character display, which may not view properly for some displays.

Example

The following is a list example of all Open Replicator sessions for array ID 0766:

```
symrcopy list -sid 0766
```

Sample output

Symmetrix ID: 000195700766

Control Device	Remote Device	Flags	Status	Done
Sym	Protected Tracks Identification	RI CDSHUTZ	CTL <=> REM	(%)
002FF	0 6001248000DF0DC47E0044E36089EFEE	.W XXXX.R.	Synchronized	100
0030D	0 60012480001A6EF92F00DFCB2408DABC	.W XXXX.R.	Synchronized	100
00311	0 60012480001A6EF92F00B7E4AC5C1408	.W XXXX.R.	Synchronized	100
00314	0 60012480001A6EF92F008F23982A6686	.W XXXX.R.	Stopped	N/A
00316	0 6001248000DF0DC47E00CCC8248867CB	.W XXXX.R.	Stopped	N/A
00317	0 6001248000DF0DC47E004D6586E5C423	.W XXXX.R.	Stopped	N/A
00318	0 6001248000DF0DC47E0004F0A43FAF17	.W XXXX.R.	Stopped	N/A
00319	0 6001248000DF0DC47E003FE69C56748B	.W XXXX.R.	Stopped	N/A
0039F	0 6001248000DF0DC47E005CC2D44B6D1C	.W XXXX.R.	Synchronized	100
003A7	0 6001248000DF0DC47E00D05F24EB2D64	.W XXXX.R.	Synchronized	100
003AF	0 6001248000DF0DC47E00D31A2E28B989	.W XXXX.R.	Stopped	N/A
003B7	0 6001248000DF0DC47E00E8719DEEB8DE	.W XXXX.R.	Synchronized	100
003BF	0 6001248000DF0DC47E006409B754B70E	.W XXXX.R.	Stopped	N/A
003C7	0 6001248000DF0DC47E004C9C0ECAD8A9	.W XXXX.R.	Stopped	N/A

Total -----
Tracks 0
MB(s) 0.0

Legend:

R: (Remote Device Vendor Identification)
S = Symmetrix, C = Clariion, . = Unknown.

I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.

Flags:

(C): X = The background copy setting is active for this pair.
. = The background copy setting is not active for this pair.

(D): X = The session is a differential copy session.
. = The session is not a differential copy session.

(S): X = The session is pushing data to the remote device(s).
. = The session is pulling data from the remote device(s).

(H): X = The session is a hot copy session.
. = The session is a cold copy session.

(U): X = The session has donor update enabled.
. = The session does not have donor update enabled.

(T): C = The session is a continuous session.
M = The session is a migration session.
R = The session is a RecoverPoint session.
S = The session is a standard ORS session.

(Z): X = The session has front-end zero detection enabled.

. = The session does not have front-end zero detection enabled.
 (*): The failed session can be reactivated.

List filtered ORS session types

Description

The list `symrcopy listcommand` provides the `-type` option that lists either standard ORS sessions or RecoverPoint sessions.

NOTE: For listing Federated Live Migration (FLM) sessions, refer to [#unique_768](#).

Examples

To list standard sessions only, enter:

```
symrcopy list -sid 6190 -type standard
```

To list RecoverPoint sessions only, enter:

```
symrcopy list -sid 90 -type recoverpoint
```

NOTE: If the `symrcopy list` command is run without the `-type` option, the default behavior is to include all sessions in the list, as shown in [List all ORS sessions on page 428](#).

Sample output

For standard session:

```
Symmetrix ID: 000000006190
Control Device          Remote Device          Flags      Status      Done
-----
Sym      Protected
Tracks      Identification      RI CDSHUTZ CTL <=> REM (%)
-----
001F8      33000      6006048000000000619053594D314640      .W X..XXS. CreateInProg N/A
001F9      33000      6006048000000000619053594D314637      .W X..XXS. CreateInProg N/A
00170      20000      6006048000000000619053594D314642      .W X..XXSX CopyInProg 50
00172      30000      6006048000000000619053594D314646      .W X..XXSX CopyInProg 75

Total
Tracks      152000
MB(s)          12062.5
```

Legend:

```
R (Remote Device Vendor Identification)
S = Symmetrix, . = Unknown.
I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
     . = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
     . = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
     . = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
     . = The session is a cold copy session.
(U): X = The session has donor update enabled.
     . = The session does not have donor update enabled.
(T): M = The session is a migration session.
     R = The session is a RecoverPoint session.
     S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
     . = The session does not have front-end zero detection enabled.
(*): The failed session can be reactivated.
```

For RecoverPoint session:

```
Symmetrix ID: 000000006190
Control Device          Remote Device          Flags      Status      Done
-----
Sym          Protected
          Tracks          Identification          RI CDSHUTZ CTL <=> REM (%)
-----
01B2          33000          6006048000000000619053594D314649  SW ...XXR. Created      75

Total
Tracks          -----
          33000
          MB (s)          2062.5
```

Query ORS session status

Description

Open Replicator session status is checked using the `symrcopy query` or `symrcopy verify` command.

The `symrcopy query` command is used to display details for remote copy sessions defined in a device file. The query command provides current status information for control/remote device pairs. If the device pair state is CopyInProg, the query command displays the percentage of copying that has completed.

Options


The following options are available for the `symrcopy query` command:

`-i`

This (interval) is used to execute the query command in repeated intervals (in seconds). The default for interval is 30 seconds if the count option is used, and 15 seconds is the minimum interval that can be specified.

`-c`

This (count) is used with the `-i` option and specifies the duration of the query intervals.

 NOTE: Estimated time to copy completion is shown in the query output when the `-c` and `-i` options are used or if the protected track count has changed since the last interval.

`-offline`

This option displays only information held in the database and does not query the array for updated session information.

`-detail`

This option displays additional device information for modified tracks, session pace, and session name. This option expands the width of the character display, which may not view properly for some displays.

`-wwn`

This option displays the full device world wide name. This option expands the width of the character display, which may not view properly for some displays.

`-summary`

This shows all possible session states and the number of sessions that are in each state. The `-wwn` and the `-detail` options cannot be used with the `-summary` option.

Examples

To query copy session status using a file, enter:

```
symrcopy query -file dev_file_1
```

To query for copy session status that will run every 30 seconds for 1 hour, enter:

```
symrcopy query -file dev_file_1 -i 30 -c 120
```

To query copy session status using a session name, enter:

```
symrcopy query -session_name rcopy_2
```

To query copy session status using the `-summary` option, enter:

```
symrcopy -file dev_file_1 query -summary
```

Sample output

This output shows a copy session status:

```
symrcopy query -file dev_file_1
```

Control Device	Remote Device	Flags	Status	Done
Protected SID:symdev SRC <=> TGT (%)	Tracks Identification		RI	CDSHUTZ
000000006190:00168	33000 6006048000000000619053594D314638	.W X..XXM.	Copied	100
000000006190:001F8	33000 6006048000000000619053594D314640	.W X..XXS.	CreateInProg	N/A
000000006190:001F9	33000 6006048000000000619053594D314637	.W X..XXS.	CreateInProg	N/A
000000006190:00170	20000 6006048000000000619053594D314642	.W X..XXSX	CopyInProg	50
000000006190:00172	30000 6006048000000000619053594D314646	.W X..XXSX	CopyInProg	75
000000006190:001B2	33000 6006048000000000619053594D314649	.W X..XXRX	CopyInProg	75
Total				
Tracks	152000			
MB(s)	12062.5			

Legend:

R (Remote Device Vendor Identification)
 S = Symmetrix, . = Unknown.
 I: (Remote Device Specification Identifier)
 D = Device Name, W = LUN WWN, World Wide Name.
 Flags:
 (C): X = The background copy setting is active for this pair.
 . = The background copy setting is not active for this pair.
 (D): X = The session is a differential copy session.
 . = The session is not a differential copy session.
 (S): X = The session is pushing data to the remote device(s).
 . = The session is pulling data from the remote device(s).
 (H): X = The session is a hot copy session.
 . = The session is a cold copy session.
 (U): X = The session has donor update enabled.
 . = The session does not have donor update enabled.
 (T): C = The session is a continuous session
 M = The session is a migration session.
 R = The session is a RecoverPoint session.
 S = The session is a standard ORS session.
 (Z): X = The session has front-end zero detection enabled.
 . = The session does not have front-end zero detection enabled.
 (*): The failed session can be reactivated.

Verify ORS session state

Syntax

To verify session status, use the following syntax:

```
symrcopy verify [-createinprog | -created | -recreateinprog |  
-recreated | -copyinprog | -copyonaccess | -copyonwrite | -copied |  
-terminateinprog | -failed | -verifyinprog | -restored | -restinprog | -precopy  
[-cycled] | -syncinprog | -synchronized | -failedback]
```

Options

-createinprog

Verifies that the copy session is in the process of being created.

-created

	Verifies that the copy session has been created.
-recreateinprog	Verifies that the copy session is in the process of being recreated (incrementally updating the targets).
-recreated	Verifies that the copy session has been recreated. Device pairs in the session have finished incrementally updating.
-copyinprog	Verifies which device pairs in the copy session are currently in the CopyInProg state (actively background copying).
-copyonaccess	Verifies which device pairs in the copy session are currently in the CopyOnAccess state (only copying the device tracks to the control device as they are being accessed on the remote device for a pull operation).
-copyonwrite	Verifies which device pairs in the copy session are currently in the CopyOnWrite state (only copying the device tracks to the remote device as they are being written to on the remote device for a push operation).
-copied	Verifies which device pairs in the copy session have finished copying data. This is the default if no option is provided.
-terminateinprog	Verifies that the copy session is in the process of terminating.
-failed	Verifies if any of the device pairs in the copy session have failed to copy.
-failedback	Verifies if a FLM session has failed back.
-verifyinprog	Verifies that all active directors for the copy session have completed copy operations.
-precopy	Verifies that the device pair is currently in the Precopy state (copying device tracks in the background without activation). Adding the <code>-cycled</code> option verifies all precopy sessions that have completed one cycle.
-restinprog	Verifies that the copy session is in the process of being restored.
-restored	Verifies that the copy session has been fully restored.
-syncinprog	Verifies which device pairs in the copy session are currently in the SynclnProg state (actively background copying).
-synchronized	Verifies which device pairs in the copy session are in the Synchronized state.

Examples

To verify copy session states for a list of devices specified in `dev_file_1`, enter:

```
symrcopy -file dev_file_1 verify
```

To verify copy session states for a list of devices specified in `dev_file_1`, enter:

```
symrcopy -file dev_file_1 verify -summary
```

Sample output

Using the `verify` command:

```
One of the device(s) in the list are in "~Copied' state.
```

Using the `verify` command with the `-summary` option:

NOTE: The one-line `verify` command output displays after the `-summary` output.

```
Device File Name : dev_file_1
RCopy Session State      Count
-----
CreateInProg             2
Created                  0
RecreateInProg           0
Recreated                0
CopyInProg               3
CopyOnAccess             0
CopyOnWrite              0
Copied                   1
SyncInProg               0
Synchronized             0
Restored                 0
RestoreInProg            0
Precopy                  0
TerminateInProg          0
Failed                   0
Stopped                  0
FailedBack               0
VerifyInProg             0
Invalid                  0
-----
Total                    6
                        Track(s)  MB(s)
-----
Total Protected          156000    12062.5

One of the device(s) in the list are in 'Copied' state.
```

NOTE: If no `verify` option is provided, then "copied" is the default state that is verified.

Terminate an ORS session

Syntax

To terminate an ORS copy session and remove it from the array, use the following syntax:

```
symrcopy terminate
```

Options

-symforce

This option is mandatory when terminating a session in the *CopyInProg*, *CopyOnAccess*, or *CopyOnWrite* state.

NOTE:

Use care when applying the `-symforce` option to terminate an active session. At termination, the receiving devices contain an incomplete data copy and should be considered invalid.

-file

Use this with `FileName` to specify the control device associated with the session. Remote devices in the file are ignored.

-session name

Use this with `SessionName` to specify the session to terminate.

-all sessions

Terminates all sessions.

Examples

To terminate a copy session associated with control device in file `dev_file_1`, enter

```
symrcopy terminate -file dev_file_1
```

To terminate a copy session named `rcopy_1`, enter:

```
symrcopy terminate -session_name rcopy_1
```

To terminate an activated copy session that has not finished copying:

```
symrcopy terminate -file dev_file_1 -symforce
```

To terminate all sessions associated with the control device in file `dev_file_1`, enter:

```
symrcopy terminate -all_sessions -symforce -file dev_file_1
```

Terminate ORS RecoverPoint session

Description

The `symrcopy` command does not support control of Open Replicator RecoverPoint sessions, however for cleanup purposes RecoverPoint sessions can be terminated and removed.

Example

To terminate a RecoverPoint session associated with the control device in the file `dev_file_1`, enter:

```
symrcopy terminate -file dev_file_1 -rp
```

Remove an ORS remote device from a session

Example

When removing a remote device from a session, it must be put in a device file. To remove a remote device associated with device in the file `dev_file_1`, enter:

```
symrcopy remove -file dev_file_1
```

Recreate an ORS session

Description

Recreating a session creates a new point-in-time copy of the data. For differential push operations only, the copy session is recreated using the `symrcopy recreate` command. The session must have been originally created with *differential copying*. Activating a recreated session begins an incremental update of the devices to copy any device tracks that changed since the last time the session actively finished copying. Up to 15 sessions are allowed for incremental updates per logical volume. Open Replicator uses the *Symmetrix Differential Data Facility* (SDDF) to set the track protection bitmaps and monitor track differences between the control and remote devices.

Options

-name

Use this option to rename a recreated session.

-precopy

Use the option to recreate a copy session to pre-copy the incremental track updates in the background without activating the session. This option is for hot push ORS operations.

Examples

To recreate and activate a copy session for incremental track updates for session name `rcopy_2`, enter:

```
symrcopy recreate -name rcopy_2 -file dev_file_3
symrcopy activate -session_name rcopy_2
```

To recreate and rename a session to `rcopy_2`, enter:

```
symrcopy recreate -name rcopy_2 -file dev_file_3 -precopy
```

NOTE:

The `-pace` option can be included in the command line to manage the speed of the replication process. See [Set the ORS session pace](#) on page 428.

Recover from failed ORS session

Description

For a *Failed* session, when no new data is indicated on the devices in the session, it is eligible to be reactivated. Use the `symrcopy query` command to check the status of a failed session. An asterisk (*) symbol next to the *Failed* status indicates the session is available for reactivation. Use the `activate` command to reactivate the failed sessions.

Examples

To query session status, enter:

 NOTE: The *Failed* sessions shown are eligible for reactivation.

```
symrcopy -file dev_file_1 query -detail
```

Sample output

```
Symmetrix ID: 000190300237
Control Device          Remote Device          Flags      Status      Done Pace Name
-----
Sym      Protected Modified
Tracks   Tracks   Identification          RI  CDSHUTZ  CTL <=> REM  (%)
-----
001F8   33000    0  006048000000000619053594D314640  SD  X.XX.S.  Failed (*)  N/A  5
N/A
001F9   33000    0  006048000000000619053594D314637  SD  X.XX.S.  Failed (*)  N/A  5  N/A
-----
Total
Tracks 66000
MB(s)      4062.4
```

Legend:

```
R (Remote Device Vendor Identification)
S = Symmetrix, . = Unknown.
I: (Remote Device Specification Identifier)
  D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
     . = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
     . = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
     . = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
     . = The session is a cold copy session.
(U): X = The session has donor update enabled.
     . = The session does not have donor update enabled.
(T): C = The session is a continuous session.
```

```
M = The session is a migration session.
R = The session is a RecoverPoint session.
S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
      . = The session does not have front-end zero detection enabled.
(*): The failed session can be reactivated.
```

Restrictions

- Reactivating failed sessions requires Enginuity version 5773 or higher.
- If there is new data on devices that are part of the session, session activation is blocked. If session activation is blocked and it is a *non-differential* session, terminate the session and re-issue the `create` and `activate` commands. This starts the copy from the beginning.
- If session activation is blocked and it is a *differential* session, recreate the session to create a new point-in-time copy.

Restoring an ORS session

Description

Restoring a session is only allowed for differential push operations. The copy session restores back to the control device by pulling back only the changed tracks from the remote device. 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.

For example, if all data is copied from the control device to the remote device(s) and then changes are made to the control device, use the `symrcopy restore` command to recover the original data from the remote device. When the command is issued, the session is recreated in restore mode and automatically activated. At the start of the restore operation, all control devices are set to *Not Ready* status. If running a hot session, control devices are returned to *Ready* status at the end of the operation (as the data begins copying). If running a cold session, the control devices remain in *Not Ready* status.

 **NOTE:** Refer to the *EMC Solutions Enabler Installation Guide* for license information required for restore functionality.

Options

- file**
Use this with `FileName` to specify the devices associated with the session.
- session name**
Use this with `SessionName` to specify the session to restore
- pace**
Sets the speed of the session. See [Set the ORS session pace](#) on page 428

Example

To restore original data from a differential push session back to the control device, enter:

```
symrcopy restore -file dev_file_3
```

Restore ORS data using donor update

Description

Differential push operations are restored using the `-donor_update` option. Using this option with the `symrcopy restore` command, a copy is maintained, on the remote device, of any new data that has been written to the control device while the session is in the process of restoring.

Options

- force**
Terminates a session that has `-donor_update` enabled.

Turns off the `-donor_update` option while session is in *RestInProg*. The session will continue to restore in its current mode without donor update.

Use to create a new session(using the same control device), recreate, or restore sessions that are in the *Restored* state when `-donor_update` is enabled.

`donor_update off -consistent`

Maintains consistency on the remote devices after the restore session is complete. Must be set before restore action.

Examples

To restore data back to the control device using the `-donor_update` option, enter:

```
symrcopy restore -file dev_file_3 -donor_update
```

NOTE: The control device will be set to not ready before the operation and then set back to its previous state after the restore has begun.

To set the `donor_update` option to off and maintain consistency on the remote devices:

```
symrcopy set donor_update off -file dev_file_3 -consistent
```

Restrictions:

- If `-donor_update` option is used, a session cannot be renamed or devices cannot be removed from a session in the *Restored* state.
- If the session is terminated, renamed, recreated, a device is removed, or another session is created using the same control device, the donor update portion of the session will automatically be deactivated and consistency on the remote devices will be lost.

Open Replicator examples

This section provides the following examples of using Open Replicator:

Example: [Perform an ORS hot pull operation](#) on page 438

Example: [Pull data online from IBM F20 to VMAX array](#) on page 442

Example: [Pushing online data from VMAX array to HDS 9960 array](#) on page 448

Example: Perform an ORS hot pull operation

This example shows how to migrate data from an older array to a VMAX array. The hardware setup consists of the control array with array ID 000187900041 (abbreviated as 41) connected to a controlling host. The remote array on the SAN is an older array (DMX). Three remote devices are each identified by their LUN WWN. Three control devices are E9, EA, and EB. The control device capacity should be equal to or larger than the remote device extents that are being copied.

NOTE: If a copy needs to be forced from a larger device to a smaller device (for example, data was initially copied to a larger device and now the same data needs to be copied back to the smaller device), this is done by including the `-force_copy` option with the `symrcopy create` command.

For online (`-hot`) copying, the control devices may be Read/Write enabled. The remote devices should not be receiving any updates from their local host.

The steps for performing a hot pull operation are:

- [Example 1 Hot pull step 1 – create device file](#) on page 439
- [Example 1 hot pull step 2 – create migration session with donor update](#) on page 439
- [Example 1 hot pull step 3 – query migration session](#) on page 439
- [Example 1 hot pull step 4 – activate migration session](#) on page 440
- [Example 1 hot pull step 5 – verify session status](#) on page 441
- [Example 1 hot pull step 6 – terminate migration session](#) on page 442

Example 1 Hot pull step 1 – create device file

The first step in an ORS copy operation is to define the control/remote device pairings in a text file. A control device or remote device is specified by either its unique LUN WWN or by a combination of the array ID and the device name (array ID:device). Enter the control devices in the left-hand column, and the remote devices in the right-hand column, as shown below in the filename `tango`:

```
vi tango
symdev=000187900041:E9   wwn=60060480000000000314353594D303737
symdev=000187900041:EA   wwn=60060480000000000314353594D303738
symdev=000187900041:EB   wwn=60060480000000000314353594D303739
```

Example 1 hot pull step 2 – create migration session with donor update

The `symrcopy create` command creates three online copy sessions so that data on the remote devices specified in file `tango` can be copied to the control devices when the copy operation is activated. The `-pull` parameter specifies that the control array is pulling the data to it. The `-hot` parameter indicates that the control array remains online during the operation. The `-name` option gives these sessions the label name `Monday`. The `-donor_update` parameter indicates that all writes to the control device from the host will also be copied to the remote device.

```
symrcopy create -name Monday -pull -hot -donor_update -file tango -noprompt
```

```
'Create' operation execution is in progress for the device list in
device file 'tango'. Please wait...
'Create' operation successfully executed for the device list
in device file 'tango'.
```

Example 1 hot pull step 3 – query migration session

The `symrcopy query` command indicates that the sessions for the control/remote device pairs in the file `tango` are in the *Created* state and are considered to be active sessions. When the control host can "see" the remote devices (in this case, a remote array), Open Replicator converts the remote device LUN WWN identifier (specified in file `tango`) to the "array ID:device" format (for example, 000000003143:0077).

```
symrcopy query -file tango
```

```
Device File Name      : tango
Control Device        Remote Device      Flags      Status      Done
-----
SID:symdev            Protected
Tracks              Identification    RI  CDSHUTZ  CTL <=> REM  (%)
-----
000187900041:000E9   138090 000000003143:00077  SD X..XXS.  Created     N/A
000187900041:000EA   138090 000000003143:00078  SD X..XXS.  Created     N/A
000187900041:000EB   138090 000000003143:00079  SD X..XXS.  Created     N/A
Total
Track(s)              414270
MB(s)                 12945.9
```

Legend:

R (Remote Device Vendor Identification)

S = Symmetrix, . = Unknown.

I: (Remote Device Specification Identifier)

D = Device Name, W = LUN WWN, World Wide Name.

Flags:

(C): X = The background copy setting is active for this pair.

. = The background copy setting is not active for this pair.

(D): X = The session is a differential copy session.

. = The session is not a differential copy session.

(S): X = The session is pushing data to the remote device(s).

. = The session is pulling data from the remote device(s).

(H): X = The session is a hot copy session.

. = The session is a cold copy session.

(U): X = The session has donor update enabled.

. = The session does not have donor update enabled.

```

(T): C = The session is a continuous session.
      M = The session is a migration session.
      R = The session is a RecoverPoint session.
      S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
      . = The session does not have front-end zero detection enabled.
(*) : The failed session can be reactivated.

```

The `symrcopy query` command activates the copy sessions for the pairings in the file `tango`. Copying from the remote array to the control array begins. At this point, migrated data on the on the control array can be accessed without waiting for the copy operation to complete.

Example 1 hot pull step 4 – activate migration session

```
symrcopy activate -file tango -noprompt
```

```

'Activate' operation execution is in progress for the device list in
device file 'tango'. Please wait...
'Activate' operation successfully executed for the device list
in device file 'tango'.

```

The `symrcopy query` command with the `-detail` option indicates that the sessions for the device pairs defined in the file `tango` are in the `CopyInProg` state and the percent (%) completion. The display also contains other details such as the pace. The default pace value of 5 provides relatively fast copy time with only a moderate impact on the application.

```
symrcopy query -file tango -detail
```

```

Device File Name      : tango
Control Device        Remote Device      Flags      Status      Done Pace Name
-----
Protected Modified
SID:symdev    Tracks  Tracks  Identification  RI  CDSHUTZ  CTL <=> REM (%)
-----
000187900041:000E9  128083  0      000000003143:0077  SD  X..XXS.  CopyInProg  7  5  Monday
000187900041:000EA  123742  0      000000003143:0078  SD  X..XXS.  CopyInProg  10  5  Monday
000187900041:000EB  127455  0      000000003143:0079  SD  X..XXS.  CopyInProg  7  5  Monday
Total
Track(s)          379280
MB(s)             11852.5
Legend:
R (Remote Device Vendor Identification)
S = Symmetrix, . = Unknown.
I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
      . = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
      . = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
      . = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
      . = The session is a cold copy session.
(U): X = The session has donor update enabled.
      . = The session does not have donor update enabled.
(T): C = The session is a continuous session.
      M = The session is a migration session.
      R = The session is a RecoverPoint session.
      S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
      . = The session does not have front-end zero detection enabled.
(*) : The failed session can be reactivated.

```

The `symrcopy query` command checks at 60-second intervals (`-i`) to verify whether the control/remote device pairs are in the Copied state.



The Pace column will report N/A for sessions where the control device is on an array running HYPERMAX OS 5977 or higher.

Example 1 hot pull step 5 – verify session status

```
symrcopy verify -i 60 -file tango
```

```
NONE of the devices are in the 'Copied' state.
NONE of the devices are in the 'Copied' state.
NOT ALL of the devices are in the 'Copied' state.
ALL of the devices are in the 'Copied' state.
```

A subsequent `symrcopy query` command indicates that the sessions for the device pairs defined in the file `tango` are now in the *Copied* state and that copying is 100% complete.

```
symrcopy query -file tango
```

```
Device File Name      : tango
Control Device                               Remote Device      Flags      Status      Done
-----
SID:symdev      Protected
Tracks          Identification    RI CDSHUTZ CTL <=> REM      (%)
-----
000187900041:000E9      0 000000003143:0077      SD X..XXS. Copied      100
000187900041:000EA      0 000000003143:0078      SD X..XXS. Copied      100
000187900041:000EB      0 000000003143:0079      SD X..XXS. Copied      100
Total
Track(s)          0
MB(s)            0.0
Legend:
R (Remote Device Vendor Identification)
S = Symmetrix, . = Unknown.
I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
. = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
. = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
. = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
. = The session is a cold copy session.
(U): X = The session has donor update enabled.
. = The session does not have donor update enabled.
(T): C = The session is a continuous session.
M = The session is a migration session.
R = The session is a RecoverPoint session.
S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
. = The session does not have front-end zero detection enabled.
(*): The failed session can be reactivated.
```

The `symrcopy list` command displays the three inactive copy sessions on the control array whose `sid` is 000187900041 (abbreviated as **41**).

```
symrcopy list -sid 41
```

```
Symmetrix ID: 000187900041
Control Device                               Remote Device      Flags      Status      Done
-----
Sym      Protected
Tracks          Identification    RI CDSHUTZ CTL <=> REM      (%)
-----
000E9      0 000000003143:0077      SD X..XXS. Copied      100
000EA      0 000000003143:0078      SD X..XXS. Copied      100
000EB      0 000000003143:0079      SD X..XXS. Copied      100
Total -----
```

```

Tracks      0
MB(s)      0.0
Legend:
R (Remote Device Vendor Identification)
S = Symmetrix, . = Unknown.
I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
     . = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
     . = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
     . = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
     . = The session is a cold copy session.
(U): X = The session has donor update enabled.
     . = The session does not have donor update enabled.
(T): C = The session is a continuous session.
     M = The session is a migration session.
     R = The session is a RecoverPoint session.
     S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
     . = The session does not have front-end zero detection enabled.
(*): The failed session can be reactivated.

```

Example 1 hot pull step 6 – terminate migration session

The `symrcopy terminate` command ends all copy sessions defined in the file `tango`.

```
symrcopy terminate -file tango -noprompt
```

```
'Terminate' operation execution is in progress for the device list in
device file 'tango'. Please wait...
'Terminate' operation successfully executed for the device list
in device file 'tango'.
```

Another `symrcopy list` command verifies that there are no longer any copy sessions on the control array.

```
symrcopy list -sid 41
```

```
Symmetrix ID: 000187900041
No Devices with RCopy sessions were found.
```

With the copy operation complete, the remote application on the remote host can be started. However, any changes to remote data at this point are not migrated to the control array unless another full Open Replicator pull operation is performed.

Example: Pull data online from IBM F20 to VMAX array

This example shows how to migrate using a hot pull from an IBM F20 array to a VMAX array. Oracle is part of the environment, as is the Veritas volume manager and file system. The example shows how to perform Open Replicator operations on the controlling host connected to the VMAX array, and how to perform operations on the remote host connected to the F20 array.

Steps for pulling online data from IBM F20 to VMAX array are:

1. [Example 3 pull data online from IBM F20 to VMAX array step 1 – identify IBM devices](#) on page 443 — gets WWNs of the IBM array devices.
2. [Example 3 pull data online from IBM F20 to VMAX array step 2 – create device file](#) on page 443
3. [Example 3 pull data online from IBM F20 to VMAX array step 3 – create migration session](#) on page 444
4. [Example 3 pull data online from IBM F20 to VMAX array step 4 – shutdown remote application running on IBM F20 devices](#) on page 444
5. [Example 3 pull data online from IBM F20 to VMAX array step 5 – activate migration session](#) on page 444
6. [Example 3 pull data online from IBM F20 to VMAX array step 6 – query migration session](#) on page 444
7. [Example 3 pull data online from IBM F20 to VMAX array step 7 – set ceiling value](#) on page 445
8. [Example 3 pull data online from IBM F20 to VMAX array step 8 – adjusting ceiling value](#) on page 446

Example 3 pull data online from IBM F20 to VMAX array step 3 – create migration session

A `symrcopy create` command from the control host now sets up the Open Replicator hot pull operation. The command creates ten online copy sessions so that data on the remote IBM devices specified in file `devfile.pull` can be copied to the control devices when the copy operation is started. The `-pull` parameter specifies that the control array is pulling the data to it. The `-hot` parameter indicates that the array application remains online during the operation. The `-name` option gives these sessions the label name IBM:

```
symrcopy create -name IBM -pull -hot -file devfile.pull -noprompt
```

```
'Create' operation execution is in progress for the device list in
device file 'devfile.pull'. Please wait...
'Create' operation successfully executed for the device list
in device file 'devfile.pull'.
```

Example 3 pull data online from IBM F20 to VMAX array step 4 – shutdown remote application running on IBM F20 devices

Although not shown here, on the remote host, shut down the remote application that uses the F20 array devices, unmount the remote file system(s), and deport volume group(s). By performing these steps after creating the Open Replicator session, this ensures that the create operation is successful and the setup is correct before incurring application down time.

Example 3 pull data online from IBM F20 to VMAX array step 5 – activate migration session

The `symrcopy activate` command activates the copy sessions for the pairings in the file `devfile.pull`. Copying from the remote IBM array to the control-side array begins. At this point the migrated data can be accessed on the control-side array. The copy operation does need to be complete:

```
symrcopy activate -file devfile.pull -noprompt
```

```
'Activate' operation execution is in progress for the device list in
device file 'devfile.pull'. Please wait...
'Activate' operation successfully executed for the device list
in device file 'devfile.pull'.
```

Immediately after a successful `activate` operation and before copy operations are complete, volume group(s) can be imported, file system(s) mounted on the control host, and the application can be run on the VMAX Family array.

Example 3 pull data online from IBM F20 to VMAX array step 6 – query migration session

The `symrcopy query` command with the `-detail` option indicates that the sessions for the device pairs defined in the file `devfile.pull` are in the *CopyInProg* state and the percent (0%) completion. The display also contains other details such as the pace. The default pace value of 5 provides relatively fast copy time with only a moderate impact on the application:

```
symrcopy query -file devfile.pull -detail
```

```
Device File Name      : devfile.pull
Control Device        Remote Device          Flags          Status
Done Pace Name
-----
-----
Modified              Protected
SID:symdev            Tracks    Tracks    Identification    RI CDSHUTZ CTL <=> REM
(%)
-----
-----
000187990125:001C4 304380    0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 0
5 IBM
000187990125:001C5 304382    0 49424D20202020203231303520202* .W X..X.S. CopyInProg 0
5 IBM
000187990125:001C6 304383    0 49424D20202020203231303520202* .W X..X.S. CopyInProg 0
5 IBM
```

```

000187990125:001C7 304384      0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 0
5 IBM
000187990125:001C8 304387      0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 0
5 IBM
000187990125:001C9 304387      0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 0
5 IBM
000187990125:001CA 304389      0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 0
5 IBM
000187990125:001CB 304390      0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 0
5 IBM
000187990125:001CC 304391      0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 0
5 IBM
000187990125:001CD 304391      0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 0
5 IBM
Total
  Track(s)      3043864
  MB(s)         95120.8

```

Legend:

```

R (Remote Device Vendor Identification)
S = Symmetrix, . = Unknown.
I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
     . = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
     . = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
     . = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
     . = The session is a cold copy session.
(U): X = The session has donor update enabled.
     . = The session does not have donor update enabled.
(T): X = The session is a migration session.
     R = The session is a RecoverPoint session.
     S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
     . = The session does not have front-end zero detection enabled.
(*): The failed session can be reactivated.

```

NOTE:

The Pace column will report N/A for sessions where the control device is on an array running HYPERMAX OS 5977 or higher.

Example 3 pull data online from IBM F20 to VMAX array step 7– set ceiling value

The ceiling value is the percentage of the bandwidth available for Open Replicator background copy transfers. This value can be set but should only be done after understanding the bandwidth being used by all other applications. There may be other applications using the same Fibre Channel director(s) as Open Replicator. Setting the Open Replicator ceiling too high for a director/port can have an adverse impact on these other applications. The ceiling settings that this example uses are for demonstration purposes only.

By default, the ceiling is undefined (as indicated by NONE in the display). The "Max" value is the estimated maximum bandwidth (MB/second) for each director/port of the control-side array. A bandwidth ceiling can be set that balances application performance against Open Replicator copy time. Because the ceiling is not set, the speed of the copy operation is currently controlled by the default pace setting (5) displayed earlier. The following `list ceiling` command to lists the current ceiling setting.

```

symqos -rcopy list ceiling -dir all -sid 25
Symmetrix ID: 000187990125
  Symmetrix Remote Copy Bandwidth Ceiling
      Max Set Actual
  Dir:P (MB) (%) (MB)
  ---- - - - -
  01C:0 130 NONE 0
  01C:1 130 NONE 0
  02C:0 130 NONE 0
  02C:1 130 NONE 0
  15C:0 130 NONE 0
  15C:1 130 NONE 0

```



```

000187990125:001CA 299774 0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 1 5
IBM
000187990125:001CB 299989 0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 1 5
IBM
000187990125:001CC 300146 0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 1 5
IBM
000187990125:001CD 301044 0 49424D2020202020323130352020202* .W X..X.S. CopyInProg 1 5
IBM
Total -----
Track(s) 2998851
MB(s) 93714.1
Copy rate : 13.0 MB/S
Estimated time to completion : 02:00:11
Legend:
R (Remote Device Vendor Identification)
S = Symmetrix, . = Unknown.
I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
. = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
. = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
. = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
. = The session is a cold copy session.
(U): X = The session has donor update enabled.
. = The session does not have donor update enabled.
(T): M = The session is a migration session.
R = The session is a RecoverPoint session.
S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
. = The session does not have front-end zero detection enabled.
(*): The failed session can be reactivated.

```

NOTE:

The Pace column will report N/A for sessions where the control device is on an array running HYPERMAX OS 5977 or higher.

Another `symqos -rcopy set ceiling` command sets a new bandwidth ceiling of 80% for director 16c, port 1, giving Open Replicator most of the possible FA bandwidth. Most likely this setting would impact any applications using director/port FA 16C:1:

```

symqos -rcopy set ceiling 80 -dir 16c -port 1 -sid 25 -noprompt
'Set Ceiling' operation execution is in progress
'Set Ceiling' operation successfully executed

```

The following `symqos -rcopy list ceiling` command displays the ceiling setting for all directors, including director 16c, port 1. Although the actual bandwidth being used (currently 37 MB/second) is not at 80% of the maximum, it may approach that value as the copy operation progresses. However, the "Actual" value is affected by the SAN and the remote storage, which may keep this value below the percentage allowed for Open Replicator. If the ceiling is never reached, then the ceiling does not affect the copy rate.

```
symqos -rcopy list ceiling -dir all -sid 25
```

```

Symmetrix ID: 000187990125
Symmetrix Remote Copy Bandwidth Ceiling
      Max Set Actual
Dir:P (MB) (%) (MB)
-----
01C:0 130 10 0
01C:1 130 10 0
02C:0 130 10 0
02C:1 130 10 0
15C:0 130 10 0
15C:1 130 10 0
16C:0 130 10 0
16C:1 130 80 37
02D:0 130 10 0
02D:1 130 10 0

```

```
16D:0 130 10 0
16D:1 130 10 0
```

The `symrcopy verify` command checks at 60-second intervals (-i) whether the control/remote device pairs are in the *Copied* state. The Open Replicator copy operation is now complete:

```
symrcopy verify -copied -file devfile.pull -i 60
```

```
NONE of the devices are in the 'Copied' state.
NONE of the devices are in the 'Copied' state.
....
ALL of the devices are in the 'Copied' state.
```

The `symrcopy terminate` command ends all copy sessions defined in the file `devfile.pull`:

```
symrcopy terminate -file devfile.pull -noprompt
```

```
'Terminate' operation execution is in progress for the device list in
device file 'devfile.pull'. Please wait...
'Terminate' operation successfully executed for the device list
in device file 'devfile.pull'.
```

With the copy operation complete, the remote application on the remote host can be restarted (if necessary). However, any changes to remote data at this point are not migrated to the VMAX array unless another full Open Replicator pull operation is performed.

Example: Pushing online data from VMAX array to HDS 9960 array

This example shows how to perform a hot push of data from a Symmetrix VMAX family array to a Hitachi HDS 9960 array. The example shows how to perform Open Replicator operations on the controlling host connected to the VMAX Family array, and how to perform operations on the remote host connected to the HDS array.

The steps for pushing online data from VMAX array to HDS 9960 array :

1. [Example 4 push online data from VMAX array to HDS 9960 array step 1 – identify HDS devices](#) on page 449 — gets WWNs of the HDS array devices.
2. [Example 4 push online data from VMAX array to HDS 9960 array step 2 – create device file](#) on page 449
3. [Example 4 push online data from VMAX array to HDS 9960 array step 3– create migration session](#) on page 449
4. [Example 4 push online data from VMAX array to HDS 9960 array step 4 – set ceiling value](#) on page 450
5. [Example 4 push online data from VMAX array to HDS 9960 array step 5 – activate migration session](#) on page 451
6. [Example 4 push online data from VMAX array to HDS 9960 array step 6– adjust ceiling value](#) on page 451

NOTE:

For arrays running HYPERMAX OS 5977, push operations are not supported.

NOTE:

Applications against remote HDS devices cannot be run at any time during the copy operation.

Example 4 push online data from VMAX array to HDS 9960 array – prerequisite tasks

Prior to running Open Replicator in this environment, the following tasks must be performed:

1. Configure the Fibre Channel switch to zone the VMAX array to the HDS array.
2. Configure the HDS array to assign devices to the VMAX array.

Example 4 push online data from VMAX array to HDS 9960 array step 1 – identify HDS devices

The example needs to identify the remote HDS devices that will receive data from the control-side array. EMC's Inquiry Utility version 7.3 (SIL version 6.0.2) can accomplish this when run on the remote host connected to the HDS array. If the Inquiry Utility is not available, use HDS tools. The following `inq` command identifies the HDS storage devices. Open Replicator needs to know the WWN of each HDS device:

```
inq -hds_wwn
Inquiry utility, Version V7.3-690 (Rev 0.38)
(SIL Version V6.0.2.0 (Edit Level 640)
Copyright (C) by EMC Corporation, all rights reserved.
For help type inq -h.
```

HDS Device	Array Serial #	WWN	Array Type
/dev/rdsk/c3t10d0s2	65535	4849544143484920523430303943424430303030	R400
/dev/rdsk/c3t10d1s2	65535	4849544143484920523430303943424430303032	R400
/dev/rdsk/c3t10d2s2	65535	4849544143484920523430303943424430303034	R400
/dev/rdsk/c3t10d3s2	65535	4849544143484920523430303943424430303036	R400
/dev/rdsk/c3t10d4s2	65535	4849544143484920523430303943424430303038	R400
/dev/rdsk/c3t10d5s2	65535	4849544143484920523430303943424430303041	R400
/dev/rdsk/c3t10d6s2	65535	4849544143484920523430303943424430303043	R400
/dev/rdsk/c3t10d7s2	65535	4849544143484920523430303943424430303045	R400

Example 4 push online data from VMAX array to HDS 9960 array step 2 – create device file

After identifying the control devices (1C4 - 1CD) that will send the data, define the control/remote device pairings in a text file. The following command uses the `vi` text editor to create a text file named `devfile.push`. The first pairing entered in this file is control device 1C4 on array 000187990125 (abbreviated as 25), paired with the HDS device whose WWN is 4849544143484920523430303943424430303030. The next control device is paired with the next remote HDS device, and so forth:

```
vi devfile.push
symdev=25:1C4 wwn=4849544143484920523430303943424430303030
symdev=25:1C5 wwn=4849544143484920523430303943424430303032
symdev=25:1C6 wwn=4849544143484920523430303943424430303034
symdev=25:1C7 wwn=4849544143484920523430303943424430303036
symdev=25:1C8 wwn=4849544143484920523430303943424430303038
symdev=25:1C9 wwn=4849544143484920523430303943424430303041
symdev=25:1CA wwn=4849544143484920523430303943424430303043
symdev=25:1CB wwn=4849544143484920523430303943424430303045
symdev=25:1CC wwn=4849544143484920523430303943424430303130
symdev=25:1CD wwn=4849544143484920523430303943424430303132
```

Example 4 push online data from VMAX array to HDS 9960 array step 3– create migration session

A `symrcopy create` command from the control host now sets up the Open Replicator hot push operation. The command creates ten online copy sessions so that data on the control devices specified in file `devfile.push` can be copied to the remote HDS devices when the copy operation is started. The `-push` parameter specifies that the control array is pushing the data to the remote array. The `-hot` parameter indicates that the array application remains online during the operation. Subsequent copying during these copy sessions will perform incremental copies, capturing only new writes to the control devices. The `-name` option gives these sessions the label name HDS:

```
symrcopy create -name HDS -push -hot -file devfile.push -noprompt
'Create' operation execution is in progress for the device list in
device file 'devfile.push'. Please wait...
'Create' operation successfully executed for the device list
in device file 'devfile.push'.
```

The `symrcopy query` command indicates that the sessions for the device pairs defined in the file `devfile.push` are in the Created state. To display more detail, include the `-detail` option. To see the full WWN identifier of each remote device, include the `-wwn` option:

```

symrcopy query -file devfile.push
Device File Name      : devfile.push
Control Device        Remote Device        Flags        Status        Done
-----
SID:symdev            Protected
Tracks              Identification        RI CDSHUTZ CTL <=> REM      (%)
-----
000187990125:001C4    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
000187990125:001C5    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
000187990125:001C6    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
000187990125:001C7    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
000187990125:001C8    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
000187990125:001C9    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
000187990125:001CA    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
000187990125:001CB    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
000187990125:001CC    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
000187990125:001CD    435150 4849544143484920523430303943424* .W XXXX.S. Created      N/A
Total
Track(s)              4351500
MB(s)                 135984
Legend:
R (Remote Device Vendor Identification)
S = Symmetrix, . = Unknown.
I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
      . = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
      . = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
      . = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
      . = The session is a cold copy session.
(U): X = The session has donor update enabled.
      . = The session does not have donor update enabled.
(T): M = The session is a migration session.
      R = The session is a RecoverPoint session.
      S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
      . = The session does not have front-end zero detection enabled.
(*): The failed session can be reactivated.

```

Example 4 push online data from VMAX array to HDS 9960 array step 4 – set ceiling value

The ceiling value is the percentage of the bandwidth available for Open Replicator background copy transfers. This value can be set but should only be done after understanding the bandwidth being used by all other applications. There may be other applications using the same Fibre Channel director(s) as Open Replicator. Setting the Open Replicator ceiling too high for a director/port can have an adverse impact on these other applications. The ceiling settings that this example uses are for demonstration purposes only.

By default, the ceiling is undefined (as indicated by NONE in the display). The "Max" value is the estimated maximum bandwidth (MB/second) for each director/port of the control-side array. A bandwidth ceiling can be set that balances application performance against Open Replicator copy time:

```

symqos -rcopy list ceiling -dir all -sid 25
Symmetrix ID: 000187990125
Symmetrix Remote Copy Bandwidth Ceiling
      Max Set Actual
Dir:P (MB) (%) (MB)
-----
01C:0 130 NONE 0
01C:1 130 NONE 0
02C:0 130 NONE 0
02C:1 130 NONE 0
15C:0 130 NONE 0
15C:1 130 NONE 0

```

```

16C:0  130  NONE  0
16C:1  130  NONE  0
02D:0  130  NONE  0
02D:1  130  NONE  0
16D:0  130  NONE  0
16D:1  130  NONE  0

```

The `symqos -rcopy set ceiling` command sets a bandwidth ceiling of 10% for all director/ports in the control-side array (`sid 25`). This means that Open Replicator's ceiling will be 10% of the estimated 130 MB/second FA bandwidth:

```

symqos -rcopy set ceiling 10 -dir all -sid 25 -noprompt
'Set Ceiling' operation execution is in progress
'Set Ceiling' operation successfully executed

```

The `symqos -rcopy list ceiling` command displays that the ceiling settings for all director/ports in the control-side array are now at 10%. Once the Open Replicator session is activated, the `ceiling` can be displayed again to show its "Actual" value. Note that the pace value (including the default) is ignored for any copy session that uses an FA where the ceiling is set:

```

symqos -rcopy list ceiling -dir all -sid 25
Symmetrix ID: 000187990125
  Symmetrix Remote Copy Bandwidth Ceiling
        Max      Set      Actual
  Dir:P  (MB)    (%)    (MB)
  -----
01C:0   130     10     0
01C:1   130     10     0
02C:0   130     10     0
02C:1   130     10     0
15C:0   130     10     0
15C:1   130     10     0
16C:0   130     10     0
16C:1   130     10     0
02D:0   130     10     0
02D:1   130     10     0
16D:0   130     10     0
16D:1   130     10     0

```

Example 4 push online data from VMAX array to HDS 9960 array step 5 – activate migration session

The `symrcopy activate` command starts the copy operation for the device pairs defined in the file `devfile.push`. Copying from the control devices to the remote devices begins. Using the `-consistent` option creates a consistent point-in-time copy:

```

symrcopy activate -file devfile.push -consistent -noprompt
'Activate' operation execution is in progress for the device list in
device file 'devfile.push'. Please wait...
'Activate' operation successfully executed for the device list
in device file 'devfile.push'.

```

Example 4 push online data from VMAX array to HDS 9960 array step 6– adjust ceiling value

The `symqos -rcopy list ceiling` command shows that the ceiling settings for all director/ports in the control-side array are now at 10%. Because the control devices are mapped only to director/port `16C:1`, copying occurs only through this director/port. Note that the "Actual" bandwidth being used by Open Replicator for this operation is 13 MB/second, which is 10% of the estimated maximum:

```

symqos -rcopy list ceiling -dir all -sid 25
Symmetrix ID: 000187990125
  Symmetrix Remote Copy Bandwidth Ceiling
        Max      Set      Actual
  Dir:P  (MB)    (%)    (MB)
  -----
01C:0   130     10     0
01C:1   130     10     0
02C:0   130     10     0

```

```

02C:1 130 10 0
15C:0 130 10 0
15C:1 130 10 0
16C:0 130 10 0
16C:1 130 10 13
02D:0 130 10 0
02D:1 130 10 0
16D:0 130 10 0
16D:1 130 10 0

```

NOTE:

The `symrcopy` query command displays the status of the copy operation at 30-second intervals:

```
symrcopy query -file devfile.push -i 30
```

Device File Name	Control Device	Remote Device	Flags	Status	Done
SID:symdev	Protected	Identification	RI CDSHUTZ CTL <=> REM		(%)
000187990125:001C4	416645	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
000187990125:001C5	416888	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
000187990125:001C6	416691	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
000187990125:001C7	416632	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
000187990125:001C8	416799	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
000187990125:001C9	417123	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
000187990125:001CA	416876	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
000187990125:001CB	417092	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
000187990125:001CC	417009	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
000187990125:001CD	416717	4849544143484920523430303943424*	.W XXXX.S.	CopyInProg	4
Total	-----				
Track(s)	4168472				
MB(s)	130265				

Legend:

```

R (Remote Device Vendor Identification)
S = Symmetrix, C = Clariion, . = Unknown.
I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
     . = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
     . = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
     . = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
     . = The session is a cold copy session.
(U): X = The session has donor update enabled.
     . = The session does not have donor update enabled.
(T): M = The session is a migration session.
     R = The session is a RecoverPoint session.
     S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
     . = The session does not have front-end zero detection enabled.
(*): The failed session can be reactivated.

```

The `symqos -rcopy set ceiling` command sets a bandwidth ceiling of 100% for all director/ports in the control-side array (sid 25), giving Open Replicator all of the possible FA bandwidth. Most likely this setting would impact any applications using director/port FA 16C:1:

```

symqos -rcopy set ceiling 100 -dir all -sid 25 -noprompt
'Set Ceiling' operation execution is in progress.
'Set Ceiling' operation successfully executed.

```

The following `symqos -rcopy` command displays the new ceiling setting for all directors, including director 16c, port 1. Although the actual bandwidth being used (currently 47 MB/second) is not at 100% of the maximum, it may approach that value as the copy operation

progresses. However, the "Actual" value is affected by the SAN and the remote storage, which may keep this value below the estimated maximum of the control-side array director/port. If the ceiling is never reached, then the ceiling does not affect the copy rate:

```

symqos -rcopy list ceiling -dir all -sid 25

Symmetrix ID: 000187990125
Symmetrix Remote Copy Bandwidth Ceiling

```

Dir:P	Max (MB)	Set (%)	Actual (MB)
01C:0	130	100	0
01C:1	130	100	0
02C:0	130	100	0
02C:1	130	100	0
15C:0	130	100	0
15C:1	130	100	0
16C:0	130	100	0
16C:1	130	100	47
02D:0	130	100	0
02D:1	130	100	0
16D:0	130	100	0
16D:1	130	100	0

Another `symrcopy query` command displays an updated status of the copy operation at 30-second intervals:

```

symrcopy query -file devfile.push -i 30
Device File Name      : devfile.push
Control Device        Remote Device        Flags      Status      Done
-----
SID:symdev           Protected
Tracks      Identification      RI CDSHUTZ CTL <=> REM      (%)
-----
000187990125:001C4  390617 4849544143484920523430303943424* .W XXXX.S. CopyInProg      10
000187990125:001C5  381968 4849544143484920523430303943424* .W XXXX.S. CopyInProg      12
000187990125:001C6  386389 4849544143484920523430303943424* .W XXXX.S. CopyInProg      11
000187990125:001C7  386463 4849544143484920523430303943424* .W XXXX.S. CopyInProg      11
000187990125:001C8  381195 4849544143484920523430303943424* .W XXXX.S. CopyInProg      12
000187990125:001C9  399176 4849544143484920523430303943424* .W XXXX.S. CopyInProg      8
000187990125:001CA  396252 4849544143484920523430303943424* .W XXXX.S. CopyInProg      8
000187990125:001CB  399542 4849544143484920523430303943424* .W XXXX.S. CopyInProg      8
000187990125:001CC  398678 4849544143484920523430303943424* .W XXXX.S. CopyInProg      8
000187990125:001CD  397450 4849544143484920523430303943424* .W XXXX.S. CopyInProg      8
Total
Track(s)          3917730
MB(s)             122429
Copy rate          : 48.3 MB/S
Estimated time to completion : 00:42:16
Legend:
R (Remote Device Vendor Identification)
S = Symmetrix, . = Unknown.
I: (Remote Device Specification Identifier)
D = Device Name, W = LUN WWN, World Wide Name.
Flags:
(C): X = The background copy setting is active for this pair.
     . = The background copy setting is not active for this pair.
(D): X = The session is a differential copy session.
     . = The session is not a differential copy session.
(S): X = The session is pushing data to the remote device(s).
     . = The session is pulling data from the remote device(s).
(H): X = The session is a hot copy session.
     . = The session is a cold copy session.
(U): X = The session has donor update enabled.
     . = The session does not have donor update enabled.
(T): M = The session is a migration session.
     R = The session is a RecoverPoint session.
     S = The session is a standard ORS session.
(Z): X = The session has front-end zero detection enabled.
     . = The session does not have front-end zero detection enabled.
(*): The failed session can be reactivated.

```

The `symrcopy verify` command checks at 60-second intervals (`-i`) to verify whether the control/remote device pairs are in the Copied state. The Open Replicator copy operation is now complete:

```
symrcopy verify -Copied -i 60 -file devfile.push
NONE of the devices are in the 'Copied' state.
NONE of the devices are in the 'Copied' state.
...
NOT ALL of the devices are in the 'Copied' state.
ALL of the devices are in the 'Copied' state.
```

Open Replicator operational restrictions and state reference

This section details which ORS operations are permissible for certain control device states, devices in various states of a replication (TimeFinder or SRDF) operation, and certain device types.

[Device control operations allowed for ORS control devices](#) on page 454

[ORS operations allowed for replication session states](#) on page 455

[ORS operations allowed for device types](#) on page 463

Device control operations allowed for ORS control devices

Table 27. [Device control operations allowed for ORS control device states with host](#) on page 454 details if a device control operation is permissible, for the ORS control device, in various states with the controlling host (such as, ready, not ready).

Table 27. Device control operations allowed for ORS control device states with host

Action	Allowed
Ready	
Online RCopy control device pushing out	Yes
Online RCopy control device pulling in	Yes
Offline RCopy control device pushing out	No
Offline RCopy control device pulling in	No
Not Ready	
Online RCopy control device pushing out	Yes
Online RCopy control device pulling in	Yes
Offline RCopy control device pushing out	N/A
Offline RCopy control device pulling in	N/A
RW Enable	
Online RCopy control device pushing out	Yes
Online RCopy control device pulling in	Yes
Offline RCopy control device pushing out	No
Offline RCopy control device pulling in	No
Write Disable	
Online RCopy control device pushing out	Yes
Online RCopy control device pulling in	Yes
Offline RCopy control device pushing out	No

Table 27. Device control operations allowed for ORS control device states with host (continued)

Offline RCopy control device pulling in	No
---	----

ORS operations allowed for replication session states

Use the state tables to determine if an ORS operation is permissible on devices in various states of a replication (TimeFinder or SRDF) session:

- TimeFinder/Mirror and TimeFinder/Clone sessions — Details if an ORS operation is supported on devices in various states of a TimeFinder session.
- TimeFinder/SnapVX sessions (HYPERMAX OS 5977) — Details if an ORS operation is supported on devices in various states of a TimeFinder/SnapVX session.
- SRDF sessions — Details if an ORS operation is supported on devices (such as R1 and R2) in various states of SRDF sessions.

TimeFinder/Mirror and TimeFinder/Clone sessions

The following tables detail which ORS (rcopy) operations allowed when the control device is in use as a source or target for TimeFinder/Mirror, or as a source or target for TimeFinder/Clone:

Table 28. Allowable rcopy operations when the control device for PUSH session is in use as a TimeFinder/Mirror SOURCE

rcopy Operation	BCV Pair State											
	Never Established	Sync in prog	Synchronized	Split in prog	Background split in progress	Split before sync	Split	Split no incremental	Restore in prog	Restored	Split before restore	Invalid
Activate	Y	Y	Y	Y	Y	Y	Y	Y		Y		
Create	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a		Y ^a		
Incremental Restore	Y	Y	Y	Y	Y	Y	Y	Y		Y	Y	Y
Recreate	Y	Y	Y	Y	Y	Y	Y	Y		Y		
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Precopy	Y	Y	Y	Y	Y	Y	Y	Y		Y		
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

a. Not supported for arrays running HYPERMAX OS 5977 or higher.

Table 29. Allowable rcopy operations when the control device for PUSH session is in use as a TimeFinder/Mirror TARGET

rcopy Operation	BCV Pair State											
	Never Established	Sync in prog	Synchronized	Split in progress	Background split in progress	Split before sync	Split	Split no incremental	Restore in prog	Restored	Split before restore	Invalid
Activate	Y						Y	Y				
Create	Y ^a						Y ^a	Y ^a				
Incremental Restore	Y					Y	Y	Y			Y	
Recreate	Y						Y	Y				
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Precopy	Y						Y	Y				
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

a. Not supported for arrays running HYPERMAX OS 5977 or higher.

Table 30. Allowable rcopy operations when the control device for PULL session is in use as a TimeFinder/Mirror SOURCE

rcopy Operation	BCV Pair State											
	Never Established	Sync in prog	Synchronized	Split in progress	Background split in progress	Split before sync	Split	Split no incremental	Restore in prog	Restored	Split before restore	Invalid
Activate	Y	Y	Y	Y	Y	Y	Y	Y		Y	Y	
Create	Y	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a		Y ^a	Y ^a	
Donor Update Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Failback	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
FrontEnd Zero Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 30. Allowable rcopy operations when the control device for PULL session is in use as a TimeFinder/Mirror SOURCE (continued)

rcopy Operation	BCV Pair State											
	Never Established	Sync in prog	Synchronized	Split in progress	Background split in progress	Split before sync	Split	Split no incremental	Restore in prog	Restored	Split before restore	Invalid
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

a. Not supported for arrays running HYPERMAX OS 5977 or higher.

Table 31. Allowable rcopy operations when the control device for PULL session is in use as a TimeFinder/Mirror TARGET

rcopy Operation	BCV Pair State											
	Never Established	Sync in prog	Synchronized	Split in progress	Background split in progress	Split before sync	Split	Split no incremental	Restore in prog	Restored	Split before restore	Invalid
Activate	Y					Y	Y	Y			Y	
Create	Y					Y ^a	Y ^a	Y ^a			Y ^a	
Donor Update Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Failback	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
FrontEnd Zero Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

a. Not supported for arrays running HYPERMAX OS 5977 or higher.

Table 32. Allowable rcopy operations when the control device for PUSH session is in use as a TimeFinder/Clone SOURCE

rcopy Operation	Pair State														
	No session	Create in prog	Created	Recreated	Precopy	Copy in progress	Copied	Copy on write	Copy on access	Split	Restore in prog	Restored	Terminate in prog	Invalid	Failed
Activate	Y ^a	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^a	Y ^{ab}	Y ^a	Y ^a	Y ^a
Create	Y ^a	Y ^{abc}	Y ^{abc}	Y ^{abc}	Y ^{abc}	Y ^{abc}	Y ^{abc}	Y ^{abc}	Y ^{abc}	Y ^{abc}	Y ^{ac}	Y ^{abc}	Y ^{ac}	Y ^{ac}	Y ^{ac}
Incremental Restore	Y ^a	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}		Y ^{abd}			
Recreate	Y ^a	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a
Remove	Y ^a	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y ^a	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SetPrecopy	Y ^a	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^{ab}	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

- a. Not supported if device is held.
- b. Not supported if device is the source of an extent-level Clone session.
- c. Not supported for arrays running HYPERMAX OS 5977 or higher.
- d. Not supported if device is the source of a Snap or VP Snap session.

Table 33. Allowable rcopy operations when the control device for PUSH session is in use as a TimeFinder/Clone TARGET

rcopy Operation	Pair State														
	No session	Create in prog	Created	Recreated	Precopy	Copy in progress	Copied	Copy on write	Copy on access	Split	Restore in prog	Restored	Terminate in prog	Invalid	Failed
Activate	Y ^a						Y ^{abd}			Y ^{abd}					
Create	Y ^a						Y ^{abcd}			Y ^{abcd}					
Incremental Restore	Y ^a						Y ^{abd}			Y ^{abd}					
Recreate	Y ^a						Y ^{abd}			Y ^{abd}					
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SetPrecopy	Y ^a						Y ^{abd}			Y ^{abd}					

Table 33. Allowable rcopy operations when the control device for PUSH session is in use as a TimeFinder/Clone TARGET (continued)

rcopy Operation	Pair State														
	No session	Create in prog	Created	Recreated	Precopy	Copy in progress	Copied	Copy on write	Copy on access	Split	Restore in prog	Restored	Terminate in prog	Invalid	Failed
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

- a. Not supported if device is held.
- b. Not supported if device is the source of an extent-level Clone session.
- c. Not supported for arrays running HYPERMAX OS 5977 or higher.
- d. Not supported if device is the source of a Snap or VP Snap session.

Table 34. Allowable rcopy operations when the control device for PULL session is in use as a TimeFinder/Clone TARGET

rcopy Operation	Pair State														
	No session	Create in prog	Created	Recreated	Precopy	Copy in progress	Copied	Copy on write	Copy on access	Split	Restore in prog	Restored	Terminate in prog	Invalid	Failed
Activate	Y ^a						Yabd			Yabd					
Create	Y ^a						Yabcd			Yabcd					
Donor Update Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Failback	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
FrontEnd Zero Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

- a. Not supported if device is held.
- b. Not supported if device is the source of an extent-level Clone session.
- c. Not supported for arrays running HYPERMAX OS 5977 or higher.
- d. Not supported if device is the source of a Snap or VP Snap session.

Table 35. Allowable rcopy operations when the control device for PULL session is in use as a TimeFinder/Clone SOURCE

rcopy Operation	Pair State														
	No session	Create in prog	Created	Recreated	Precopy	Copy in progress	Copied	Copy on write	Copy on access	Split	Restore in prog	Restored	Terminate in prog	Invalid	Failed
Activate	Y ^a	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}	Y ^{abd}		Y ^{abd}		Y ^{abcd}			
Create	Y ^a	Y ^{abcd}	Y ^{abcd}	Y ^{abcd}	Y ^{abcd}	Y ^{abcd}	Y ^{abcd}	Y ^{abcd}	Y ^{abcd}	Y ^{abcd}		Y ^{abd}			
Donor Update Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Failback	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
FrontEnd Zero Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

- a. Not supported if device is held.
- b. Not supported if device is the source of an extent-level Clone session.
- c. Not supported for arrays running HYPERMAX OS 5977 or higher.
- d. Not supported if device is the source of a Snap or VP Snap session.

Snap VX sessions

Only the `rcopy terminate` command is allowed when the `rcopy` control device for a PUSH or PULL session is also in use as a Snap VX source device. The SnapVX operation can only be in the following states:

- No session
- Establish in progress
- Established
- Restore in prog
- Restored
- Terminate in prog
- Failed

Only the `rcopy terminate` command is allowed when the `rcopy` control device for a PUSH or PULL session is also in use as a Snap VX target device. The SnapVX operation can only be in the following states:

- No session
- Failed
- Link copy in progress
- Link Copied

SRDF sessions

The following tables detail which ORS (`rcopy`) operations are allowed when the `rcopy` control device is in use as a SRDF R1 or R2 mirror:

Table 36. Allowable rcopy operations when the control device for PUSH session is in use as an SRDF R1 mirror

rcopy Operation	SRDF State											
	Sync in prog	Synchronized	Split	Suspended	Failed over	Partitioned1	Partitioned2	R1 updated	R1 updinprog	Invalid	Consistent	TransmitIdle
Activate	Y ^b	Y	Y	Y		Y	Y				Y ^b	Y ^b
Create	Y ^{ab}	Y ^a	Y ^a	Y ^a		Y ^a	Y ^a				Y ^{ab}	Y ^{ab}
Incremental Restore	Y ^{bc}	Y	Y	Y		Y	Y				Y ^{bc}	
Recreate	Y ^b	Y	Y	Y		Y	Y				Y ^b	Y ^b
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Precopy	Y ^b	Y	Y	Y		Y	Y				Y ^b	Y ^b
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

- a. Not supported for arrays running HYPERMAX OS 5977 or higher.
- b. Must have no local invalids on R1 or remote invalids on R2.
- c. Array must be running Enginuity 5876 or higher.

Table 37. Allowable rcopy operations when the control device for PUSH session is in use as an SRDF R2 mirror

rcopy Operation	SRDF State											
	Sync in prog	Synchronized	Split	Suspended	Failed over	Partitioned1	Partitioned2	R1 updated	R1 updinprog	Invalid	Consistent	TransmitIdle
Activate			Y	Y	Y	Y	Y	Y	Y			
Create			Y ^a	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a	Y ^a			
Incremental Restore			Y	Y	Y	Y	Y	Y ^b				
Recreate			Y	Y	Y	Y	Y	Y	Y			
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Precopy			Y	Y	Y	Y	Y	Y	Y			
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

- a. Not supported for arrays running HYPERMAX OS 5977 or higher.
- b. R2 must be write enabled.

Table 38. Allowable rcopy operations when the control device for PULL session is in use as an SRDF R1 mirror

rcopy Operation	SRDF State											
	Sync in prog	Synchronized	Split	Suspended	Failed over	Partitioned1	Partitioned2	R1 updated	R1 updinprog	Invalid	Consistent	Transmitidle
Activate	γabcd	γabcd	Y	Y		Y	Y				γabcd	
Create	γabcde	γabcde	γe	γe		γe	γe				γbcde	
Donor Update Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Failback	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Frontend Zero Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

- a. Array must be running 5876 or higher.
- b. Must have no local invalids on R1 or remote invalids on R2.
- c. Frontend zero must be in Off state.
- d. Force flag must be set.
- e. Not supported for arrays running HYPERMAX OS 5977 or higher.

Table 39. Allowable rcopy operations when the control device for PULL session is in use as an SRDF R2 mirror

rcopy Operation	SRDF State											
	Sync in prog	Synchronized	Split	Suspended	Failed over	Partitioned1	Partitioned2	R1 updated	R1 updinprog	Invalid	Consistent	Transmitidle
Activate			Y		Y	Y	Y	γbc				
Create			γa		γa	γa	γa	γabc				
Frontend Zero Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Failback	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Donor Update Off	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Nocopy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Copy	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Set Pace	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Terminate	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Remove	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Rename	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

- a. Not supported for arrays running HYPERMAX OS 5977 or higher.
- b. R2 must be write enabled.

Table 39. Allowable rcopy operations when the control device for PULL session is in use as an SRDF R2 mirror (continued)

c. Frontend zero detect must be in Off state.

ORS operations allowed for device types

The following table details if an ORS operation is permissible for certain device types.

NOTE:

Open Replicator fully supports copy operations for thin devices.

Table 40. ORS operations allowed by device type

Action	Allowed
Rcopy Create push where local device type is:	
Gatekeeper	Yes, as long as it is not the gatekeeper for the syscall.
WORM	No
CKD_3380 and CKD_3390	No
AS400	No
Rcopy Create pull, where local device type is:	
Gatekeeper	Yes, as long as it is not the gatekeeper for the syscall.
WORM	No
CKD_3380 and CKD_3390	No
AS400	Yes on DMX arrays running Enginuity 5773.150 or higher, and VMAX Family arrays running Enginuity 5876.
Rcopy Create push or pull, where local device type is:	
Virtual device (VDEV)	Yes, as long as the following conditional are met: <ul style="list-style-type: none"> · Session is a cold push ONLY · VMAX Family arrays are running Enginuity 5876.
SFS device	No
STAR	Yes
Unconfigured device	No
Meta member	No

Security related settings and procedures

This section includes settings and procedures related to Solutions Enabler security that is not included in the chapters covering [host-based](#) or [user based](#) control access.

Chapters include:

Topics:

- [Log files and settings](#)
- [Port Usage](#)
- [Lockbox](#)

Log files and settings

This chapter describes the various log files that record array-based software and hardware changes, software-based SYMAPI conditions and errors, and daemon conditions and errors.

Topics:

- [Secure audit log](#)
- [SYMAPI log files](#)
- [Daemon log files](#)

Secure audit log

Information from the secure audit log is retrieved using the Solutions Enabler `symaudit` command. The audit log records configuration changes, security alarms, service operations, and security-relevant actions on each array. Records are written to the audit log by:

- Solutions Enabler
- Software running on the service processor
- The HYPERMAX OS environment.

The audit log is maintained on the storage array itself. Event contents in the audit log cannot be altered. User access is restricted to the Auditor role that allows a user to view, but not modify, the log.

Solutions Enabler associates a unique Activity ID with each command that is executed. This Activity ID is stored within audit log records, and can be used to correlate records that belong together, such as records generated during execution of the same command.

You can configure the Solutions Enabler event daemon, `storevntd`, to automatically stream audit entries as they appear to an external log service, such as syslog, Simple Network Management Protocol (SNMP), or the Windows Event Service.

For more detail on the audit log refer to [Common audit log](#) on page 370.

SYMAPI log files

One log file is created per day, using the dated format to record SYMAPI errors and other significant conditions. A new log file is started every day on the first write after 12:00 a.m.

For more details about the log files refer to [Events and logs overview](#) on page 366

Daemon log files

Each Solutions Enabler daemon has two log files to record daemon errors and other significant conditions.

Logging alternates between the two files, switching to the other file each time the maximum size specified by the daemon's `LOGFILE_SIZE` parameter is reached. Each daemon writes to the `.log0` file until its size exceeds that specified in the `LOGFILE_SIZE` option, at which point it switches to the `.log1` file. It switches back to `.log0` under the same conditions.

For more details on daemon log files refer to [Daemon log files](#) on page 373.

Port Usage

This chapter describes the ports used by the Solutions Enabler server and the event daemon, and how to modify port settings.

Topics:

- [Port usage](#)

Port usage

This section describes the ports Solutions Enabler uses to communicate between server and client hosts.

If a firewall or network address translator is present, these ports must be open. Typically, this is a firewall between the Solutions Enabler client and the server hosts.

Server ports

In client/server mode, the Solutions Enabler server (storsrvd daemon) listens by default at TCP/IP port 2707 for client connections.

You can configure a port by adding an entry to `<SYMAPI_HOME>/config/daemon_options` file. If you change the default port at the server, you must modify the `<SYMAPI_HOME>/config/netcnfg` configuration file at client hosts to reflect the use of the non-default port.

To change the server port the server must be down. To use a different port, specify it in the `daemon_options` file, then restart the `storsrvd` daemon.

Event daemon ports

When using the asynchronous events in client/server mode, the event daemon at the client host listens at a TCP/IP port for events being forwarded from the event daemon at the server. By default, the client event daemon asks the operating system to pick an unused port for it to use.

You can configure a specific port to use by adding an entry to the `<SYMAPI_HOME>/config /daemon_options` file on the client host. The event daemon uses the following ports by default:

Port	Description
Dynamically assigned 1024 - 65535	In client/server mode, the event daemon (storevntd) on a client host listens on this port for asynchronous events sent to it from a server host. By default, this is picked at random by the client host event daemon.
514	Port the server listens on for events.
162	Port the application listens on for traps.

Configure Solutions Enabler server and event daemon ports

Server port

In client/server mode, the Solutions Enabler server (the `storsrvd` daemon) listens by default at TCP/IP port 2707 for client connections. To add an entry and configure a port for the server, add the following entry to the `daemon_options` file:

```
storsrvd:port = nnnn
```

Event daemon

When using the asynchronous events in client/server mode, the event daemon at the client host listens at a TCP/IP port for events being forwarded from the event daemon at the server. By default, the client event daemon asks the operating system to pick an unused port for it to use. To configure a port for the event daemon, add the following entry to the `daemon_options` file:

```
storevntd:event_listen_port = nnnn
```

Lockbox

This chapter describes the Solutions Enabler encrypted lockbox and the Stable System Values (SSVs), and the procedures to change the lockbox password and SSVs.

Topics:

- [Lockbox](#)

Lockbox

Solutions Enabler uses a Lockbox to store and protect sensitive information. The Lockbox is associated with a particular host. This association prevents the Lockbox from being copied to a second host and used to obtain access.

The Lockbox is created at installation. During installation, the installer is prompted to provide a password for the Lockbox. If no password is provided at installation, a default password: *nodename@SELockbox1* is generated and stored in the Lockbox along with Stable System values (SSVs, a fingerprint that uniquely identifies the host system). The *host_name* is the same value as returned by the *hostname* command.

Stable System Values (SSVs)

Stable System values (SSVs) validate access to the Lockbox. When data is written to or retrieved from the Lockbox, the SSVs in the Lockbox are compared against the SSVs generated from the host. If the SSVs match, the operation is permitted. If the SSVs do not match, the operation fails.

When Solutions Enabler is upgraded, product information in the existing Lockbox is automatically copied into the Lockbox when the Lockbox is first accessed.

When any of the following occur, the host fingerprint may no longer match, and the SSVs inside of the Lockbox must be reset:

- The host is upgraded (either hardware or software)
- The Lockbox file is moved to another host
- User clones a virtual machine

 **NOTE:** To improve security, change the Lockbox password after resetting the SSVs.

Verify Stable System Values

Description

If a host is upgraded or reconfigured host SSVs can change. Use the `symcfg -lockbox verify -ssv` command to verify the current SSVs stored in SE lockbox against the current host SSVs.

Examples

An SSV match success:

```
symcfg verify -lockbox -ssv
```

```
The Lockbox SSVs are consistent with the host System Stable Values.
```

An SSV match failure:

```
symcfg verify -lockbox -ssv
```

```
The host System Stable Values do not match the current system configuration
```

Lockbox passwords

If you create the Lockbox using the default password during installation, change the password immediately after installation to best protect the contents in the Lockbox.

For maximum security, select a password that is hard to guess. It is very important to remember the password.

 **WARNING: Loss of this password can lead to situations where the data stored in the Lockbox is unrecoverable.**

Passwords must meet the following requirements:

- 8 - 256 characters in length
- Include at least one numeric character
- Include at least one uppercase and one lowercase character
- Include at least one of the following non-alphanumeric characters: ! @ # % &

Lockbox passwords may include any character that can be typed in from US standard keyboard.

- The new password must not be the same as the previous password.

Default Lockbox password

When you install Solutions Enabler, you are asked whether you want to use the default password for the Lockbox. If you choose to use the default, the installation process establishes the default Lockbox password in the following format:


`nodename@SELockbox1`

where: *nodename* is the hostname of the computer on which you are installing.

Operating systems have different methods of determining the node name:

- Unix: The installation program uses the `hostname` command to determine the node name. Normally, the node name is set in the `/etc/hosts` file.
- Windows: The value of the `COMPUTERNAME` system environment variable, converted to lower case.
- z/OS: The `gethostname()` function is used to get the node name of the machine.

If the value of *nodename* is stored in upper case letters, it is converted to lower case for the default password.

 **NOTE: It is strongly recommended that you change the default password. If you allow the installation program to use the default password, note it for future use. You will need the password if you need to reset the Lockbox Stable System values or generate certificates for client/server operations.**

Password and SSV management

SSVs are platform-dependent. Changes to hardware or software on a host may require you to reset the SSVs stored in the Lockbox.

Lockbox administrative interactions include:

- Changing the password used to protect the Lockbox.
- Resetting the saved SSVs in the Lockbox after attributes on the host change, making the Lockbox inaccessible to user-initiated SYMAPI and SYMCLI calls.

 **NOTE: You must restart all Solutions Enabler daemons after changing or resetting the Lockbox password.**

Two `symcfg` commands allow administrative interactions with the Lockbox:

```
symcfg -lockbox [-password <Password>]
      reset -ssv
      setpw [-new_password <NewPassword>]
```

 **NOTE: Both commands require the existing password.**

Index

A

- Abort configuration session [173](#)
- access control backup file [69](#)
- access control configuration [54](#)
- access control database [49](#)
- access control operations [53](#)
- access control permissions [61](#)
- access control strategies [66](#)
- access group permissions [60](#)
- access groups [55](#)
- ACL setup [54](#)
- activate edisk [275](#)
- activate ORS session [426](#)
- add child storage groups to storage group [99](#)
- add devices to access pool [59](#)
- add devices to device group [84](#)
- add devices to storage group [97](#), [98](#)
- add devices to storage groups [285](#)
- add devices to target list [85](#)
- add director [202](#)
- add edisk [269](#)
- add host access id to group [55](#)
- add initiators [291](#)
- add initiators to initiator group [290](#)
- add initiators using input file [291](#)
- add IP route [224](#)
- add ports to port groups [288](#)
- add resource to storage container [240](#)
- add storage group Snapshot policy [101](#)
- add ungrouped devices [86](#)
- add user access id to group [56](#)
- add virtual devices to device group [85](#)
- adding edisk restrictions [270](#)
- advise, reservations [192](#)
- allocate space thin devices [71](#)
- array external locks [323](#)
- array-level data [309](#)
- attach ip interface to iscsi target [222](#), [223](#)
- audience [20](#)
- audit log for single array [370](#)
- audit log for specified time period [371](#)
- authentication types [39](#)
- authorization rules [38](#)
- auto-provisioning groups [277](#)

B

- background operations on thin devices [73](#)
- backup and restore masking view [284](#)
- backup the authorization database [46](#)
- Bring front-end directors online [323](#)
- Bring RA directors online [322](#)

C

- cascaded groups [104](#)
- cascaded initiator groups [291](#)
- cascaded storage group details
 - grouping [113](#)

- cascaded storage groups [105](#)
- ceiling value setting [427](#)
- change device state [343](#)
- change device state using file [343](#)
- change host I/O limits [109](#)
- change storage group workload type [103](#)
- changing database file [34](#)
- check free physical disk space [174](#)
- CKD devices
 - assigning PAV aliases [195](#)
- Clerra devices [87](#)
- cli commands for protocol endpoint [242](#)
- command line help [25](#)
- comments [20](#)
- compatibility [28](#)
- compatibility restrictions [27](#)
- composite groups [121](#)
- compression [149](#)
- compression control operations [149](#)
- compression reporting [151](#), [153](#)
- configuration change [170](#)
- Configuration change guidelines [174](#)
- configuration change session rules [171](#)
- configuration data overview [305](#)
- configuration data synch [35](#)
- configure server and even daemon ports [466](#), [468](#)
- connectivity authorization [32](#)
- control operations [250](#), [389](#)
- conventions for publication [20](#)
- convert a thin device [187](#)
- convert cascaded group to storage group [106](#)
- convert devices [185](#)
- convert storage group to cascaded group [105](#)
- copy all devices between storage groups [119](#)
- copy device between storage groups [119](#)
- copy devices between device groups [92](#)
- copy devices from device group to storage group [93](#)
- copy groups and masking views [284](#)
- copy port groups [289](#)
- create access group [55](#)
- create access pool [58](#)
- Create alternate ACCID [67](#)
- create and manage ACEs [60](#)
- create auto-provisioning session [278](#)
- create device group [84](#)
- create device restrictions [182](#)
- create devices [178](#), [180](#)
- Create empty device group [84](#)
- create external disk group [268](#)
- create groups and views [279](#)
- create initiator groups [281](#)
- create IP interface [213](#)
- Create iSCSI target [216](#), [218](#)
- create journal devices [252](#)
- create masking view [278](#), [281](#)
- create or modify access control data [53](#)
- create ORS copy session [424](#)
- create port groups [280](#)
- create protocol endpoint [242](#)
- create repository [258](#)

- create storage container [239](#)
- create storage groups [95](#), [280](#)
- customize SL RT multiplier [178](#)

D

- D@RE [312](#)
- data compression [149–155](#)
- data migration [379](#), [380](#), [382](#), [385–399](#), [401–404](#), [406](#), [407](#), [410](#), [411](#), [413](#)
- database access modes [34](#)
- database configuration information [35](#)
- database file location [33](#)
- database file lock [34](#)
- database location client/server mode [34](#)
- default access control configuration [66](#)
- delete access group [58](#)
- delete access pool [59](#)
- delete CHAP [296](#)
- delete device group [94](#)
- delete external disk group [269](#)
- delete initiator groups [292](#)
- delete IP interface [215](#)
- delete iscsi target [220](#)
- delete journal devices [254](#), [255](#)
- delete masking view [283](#)
- delete port groups [289](#)
- delete storage container [240](#)
- delete storage group [121](#)
- delete storage groups [287](#)
- device
 - identifiers [189](#)
- device attribute values [189](#)
- device control operations allowed [454](#)
- Device conversion restrictions [186](#)
- device emulation [344](#)
- device external locks [358](#)
- device group control operations [95](#)
- device group definition [144](#)
- device group names [81](#)
- device group types [82](#)
- device groups
 - adding devices to [123](#)
 - creating [123](#)
 - exporting [115](#)
 - RDF1 [123](#)
 - RDF2 [123](#)
- device information [422](#)
- device list by array id [306](#)
- device list capacity output [358](#)
- device list with pddevfile format [307](#)
- device list with wwn [307](#)
- device list without capacity [306](#)
- device lists [82](#)
- device masking [277](#)
- device reservations [191](#)
- device restrictions [82](#)
- device types [337](#)
- device-level data [336](#)
- devices
 - deleting [190](#)
 - mapping [194–196](#)
- director types [318](#)
- disable CHAP [296](#)

- disable donor update [425](#)
- Disable environment variables [26](#)
- Disable front-end zero detection [425](#)
- discover host HBAs [279](#)
- discover overview [30](#)
- discovery syntax [30](#)
- disk space
 - configuring into devices [184](#)
- disk-level data [358](#)
- display auto-provisioning information [297](#)
- display CHAP [296](#)
- display devices with no assignments [302](#)
- Display environment variables [25](#)
- Display full group names [26](#)
- display masking views [298](#)
- display virtual disks [33](#)
- donor update [425](#)

E

- eDisk encapsulated [268](#)
- eDisks [268](#)
- edit and manage access pools [59](#)
- enable CHAP [295](#)
- enable CHAP on iscsi initiator [295](#)
- enable donor update [425](#)
- Enable environment variables [26](#)
- Enable front-end zero detection [424](#)
- Enable user authorization [43](#)
- enforce, reservations [192](#)
- enforcement [43](#), [44](#)
- Engenuity support [19](#)
- environment expand [253](#)
- environment expand actions [253](#)
- environment setup [251](#)
- environment setup actions [250](#)
- Environmental data [329](#)
- establish administrator [68](#)
- event daemon [374](#)
- events and logs overview [366](#), [370](#)
- export device list
 - grouping [115](#)
- export device lists [89](#)
- export groups to files [115](#)
- export ORS device list [423](#)
- external spindle devices [340](#)
- external spindle information [365](#)

F

- FAST HYPERMAX [156](#)
- FAST reporting [158](#)
- FAST.X overview [266](#)
- FAST.X Solutions Enabler support [267](#)
- FCID lockdown [289](#), [290](#)
- fibre protocols [206](#)
- filter by DA, interface, disk, hypervolume [357](#)
- filter by devices in data migration [357](#)
- filter by DIF1 [357](#)
- filter by director [356](#)
- filter by director port mapping [356](#)
- filter by metadvice [356](#)
- filter by RAID information [357](#)
- filter by SRDF devices [356](#)

- filter by storage group [357](#)
- filter by technology type [357](#)
- forcible generation of alternate access ID [51](#)
- free all written tracks [73](#)
- front-end zero detection [424](#)

G

- GCM [381](#)
- General absolute control [68](#)
- GNS
 - active group list [137](#)
 - composite groups [137](#)
 - daemon [141](#)
 - output [144](#)
 - SRDF option [137](#)
- GNS daemon
 - controls [139](#)
- GNS, daemon options
 - remote mirrors [143](#)
- grant base permissions [63](#)
- grant bcv and sdr permissions [64](#)
- grant permissions to unknown user [64](#)
- grant srdf permissions [64](#)
- group management iscsi target [230](#)
- group name services [81](#)

H

- host
 - updating [199](#)
- host access control overview [49](#)
- host access id [65](#)
- host configuration database [31](#)
- host configuration database file [33](#)
- host connections [317](#)
- host connections to arrays [315](#), [316](#)
- Host ID [40–42](#)
- hot pull activate migration session [440](#)
- Hot pull create device file [439](#)
- hot pull create migration session with donor update [439](#)
- hot pull operation [438](#)
- hot pull query migration session [439](#)
- hot pull terminate migration session [442](#)
- hot pull verify session status [441](#)

I

- import device lists [90](#)
- import or export device lists [89](#)
- inhibit database synchronization [36](#)
- initial setup [68](#)
- initiator group flags [292](#)
- iscsi configuration example [226](#)
- iscsi target reporting [230](#)
- iscsi targets [212](#)

L

- limit access to access pools [58](#)
- Limit access to UnknwGrp group [67](#)
- list access control [65](#)
- list actual disk capacity [360](#)

- list all array locks [323](#)
- list all device groups [88](#)
- list all LRU groups [325](#)
- list allocations across multiple pools [351](#)
- list and show storage groups [287](#)
- list array component environmental data [330](#)
- list array environmental data [329](#)
- list array IP addresses [227](#)
- list array IP routes [229](#)
- list array iscsi targets [228](#)
- list arrays [310](#)
- list changed devices [353](#)
- list CHAP [234](#)
- list clusters [260](#)
- list cu images [326](#), [327](#)
- list cu splits [328](#)
- List data and save devices [348](#)
- list data encryption audit log [313](#)
- list data encryption status [313](#)
- list device assignments [233](#)
- list device emulations [345](#)
- list device encryption [344](#)
- list device identifiers [309](#)
- list device information [233](#), [337](#)
- list device mapping information [309](#)
- list device pools [331](#)
- list devices in device group [87](#)
- list devices mapped to directors [321](#)
- list director data [318](#)
- list director details [320](#)
- list director port data [320](#)
- list director type [319](#)
- list disk gaps [360](#)
- list disk group summary [360](#)
- list disk group summary by engine [361](#)
- list disk groups [360](#)
- list disk information [359](#)
- list engine patches [329](#)
- list environmental data example [330](#)
- list environmental data for service state [330](#)
- list events [369](#)
- list external locks [358](#)
- list external spindle state [272](#)
- list external spindles [340](#)
- list feature registrations and usage data [336](#)
- list filtered ORS session types [430](#)
- list hba [233](#)
- list hba information [308](#)
- list host HBAs [279](#)
- list initiator group devices [302](#)
- list lock details [324](#)
- list lock number details [323](#)
- list logins [234](#)
- list memory board information [326](#)
- list next available device address [321](#)
- list ORS sessions [428](#)
- list PE and VVol devices [246](#)
- list pinned devices [355](#)
- list port flags [319](#)
- list RAID group information [341](#)
- list SAS drives [359](#)
- list Service Level [159](#)
- list Snapshot policy [332](#)
- list spare physical disks [362](#)
- list specific cluster [261](#), [263](#)

- list spindle ID information [339](#)
- list spindle information [363](#)
- list spindle information for disk group [364](#)
- list spindle information for physical disk [363](#)
- list SRP [161](#)
- list SRP details [163](#)
- list SRP for thin devices [168](#)
- list storage containers for VVols [244](#)
- list storage group demand [235](#)
- list storage groups [110](#)
- List storage groups [160](#)
- list symapi services [325](#)
- list system resource demand [311](#)
- list thin bcv devices [350](#)
- list thin device tier allocations [349](#)
- list thin devices [348](#)
- list thin devices in pool [350](#)
- list user roles [47](#)
- list virtualized devices [354](#)
- locked access control session [66](#)
- log file configuration options [366](#)
- logical device names [86](#)

M

- mainframe cu image [326](#)
- manage access group [56](#)
- manage default access [67](#)
- manage initiator groups [290](#)
- manage masking view [283](#)
- manage port groups [287](#)
- manage storage groups [285](#)
- mapping FBA devices to front-end ports [194–196](#)
- merging storage groups [99](#)
- migrated devices [353](#)
- modified CLIs for vvolts [239](#)
- modified reporting command for 5977 [298](#)
- modify device capacity [183](#)
- modify IP interface [214](#)
- modify resource for storage container [241](#)
- modify role mappings [44](#)
- modify SRP [157](#)
- modify storage container description [240](#)
- modify storage groups [100](#)
- monitor authorization [46](#)
- monitor events [369](#)
- monitor thin pools [74](#)
- move access id to group [57](#)
- move all devices between device groups [92](#)
- move all devices between storage groups [118](#)
- move device between storage groups [117](#)
- move devices between device groups [91](#)
- multi session consistency [124](#)
- multiple cores and prots [319](#)
- multiple user roles [38](#)
- multisession consistency (MSC) [123](#)
- MVS subsystem IDs [184](#)

N

- name groups and views [283](#)
- naming conventions [248](#)
- netcnfg file [325](#)
- Non-Disruptive Migration [250](#), [379](#), [382](#), [389](#)

- Non-Disruptive Migration actions and states [386](#)
- Non-Disruptive Migration and other cl commands [388](#)
- Non-Disruptive Migration and other replication sessions [387](#)
- Non-Disruptive Migration cancel action [397](#)
- Non-Disruptive Migration cancel migration [407](#), [410](#)
- Non-Disruptive Migration commit action [396](#)
- Non-Disruptive Migration create action [392](#), [393](#)
- Non-Disruptive Migration create migration pathway [389](#)
- Non-Disruptive Migration cutover action [394](#), [395](#)
- Non-Disruptive Migration environment configuration [389](#)
- Non-Disruptive Migration environmental requirements [380](#)
- Non-Disruptive Migration list action [402](#)
- Non-Disruptive Migration list action use case [403](#), [411](#), [413](#)
- Non-Disruptive Migration pre-migration rules [381](#)
- Non-Disruptive Migration recover action [399](#)
- Non-Disruptive Migration remove environment [401](#)
- Non-Disruptive Migration session creation [391](#)
- Non-Disruptive Migration setup environment [390](#)
- Non-Disruptive Migration states [385](#)
- Non-Disruptive Migration sync start action [398](#)
- Non-Disruptive Migration sync stop action [398](#)
- Non-Disruptive Migration synchronize devices [398](#)
- Non-Disruptive Migration typical migration process [404](#), [406](#)
- Non-Disruptive Migration validate create action [391](#)
- Non-Disruptive Migration verify environment [389](#)
- nonpooled device permissions [68](#)

O

- online offline modes [34](#)
- Open Replicator [416](#)
- ORS [416–418](#)
- ORS and host [416](#)
- ORS background copy [427](#)
- ORS background copying [427](#)
- ORS ceiling [428](#)
- ORS copying limitation [418](#)
- ORS device file [423](#)
- ORS device guidelines [418](#)
- ORS hot pull example [420](#)
- ORS operation [420–422](#)
- ORS operational rules and limitations [418](#)
- ORS operations allowed for device types [463](#)
- ORS SAN setup [419](#)
- ORS state tables for replication [455](#)
- output
 - command line time-saving [27](#)
 - compatibility mode [28](#)
 - truncated output display [27](#)

P

- PAV alias addresses
 - adding to mapped devices [195](#)
- point-in-time copy [427](#)
- port characteristics
 - copying from one port to another [206](#)
 - setting [206](#)
- port flags [206](#)
- port group information [289](#)
- post drive replacement audit log [314](#)
- preface [20](#)
- Preset names and IDs [27](#)
- preview and commit user authorization [45](#)

- prior scripts
 - scripts [28](#)
- protect journal devices [256](#)
- protect storage group devices [255](#)
- provisioning limits [277](#)
- pull data online from IBM F20 to VMAX [442](#)
- pull data online from IBM F20 to VMAX activate migration session [444](#)
- pull data online from IBM F20 to VMAX create device file [443](#)
- pull data online from IBM F20 to VMAX create migration session [444](#)
- pull data online from IBM F20 to VMAX ID IBM devices [443](#)
- pull data online from IBM F20 to VMAX prerequisite tasks [443](#)
- pull data online from IBM F20 to VMAX query migration session [444](#)
- pull data online from IBM F20 to VMAX set ceiling value [445](#)
- pull data online from IBM F20 to VMAX shutdown remote application [444](#)
- push online data from VMAX array to HDS 9960 array [448](#)
- push online data from VMAX array to HDS 9960 array- activate migration session [451](#)
- push online data from VMAX array to HDS 9960 array- adjust ceiling value [451](#)
- push online data from VMAX array to HDS 9960 array- create device file [449](#)
- push online data from VMAX array to HDS 9960 array- create migration session [449](#)
- push online data from VMAX array to HDS 9960 array- ID HDS devices [449](#)
- push online data from VMAX array to HDS 9960 array- prereqs [448](#)
- push online data from VMAX array to HDS 9960 array- set ceiling value [450](#)

R

- rcopy concepts [417](#)
- RDF consistency group [123](#)
- reclaim allocations on thin devices [72](#)
- recover from failed session [436](#)
- recoverpoint [248](#)
- RecoverPoint integration [248](#), [250–261](#), [263](#), [264](#)
- recreate ors session
 - ORS [435](#)
- related documentation [20](#)
- release array lock [324](#)
- release device external locks [358](#)
- release locked access control session [66](#)
- remove access ID [57](#)
- remove all ACEs from access group [57](#)
- remove device from device group [93](#)
- Remove devices from access pool [59](#)
- remove devices from device group [94](#)
- remove devices from storage group [106](#)
- remove devices from storage groups [286](#)
- remove director [203](#)
- remove edisk [271](#)
- remove environment [259](#)
- remove IP route [225](#)
- remove permissions from access group [64](#)
- remove ports from port groups [288](#)
- remove remote device from session [435](#)
- remove resource from storage container [241](#)
- remove rp environment [259](#)
- remove SL and SRP from storage group [103](#)

- remove SL from storage group [102](#)
- remove SRP from storage group [103](#)
- remove storage group workload [103](#)
- Remove UnknwGrp [67](#)
- rename device groups [90](#)
- rename hba [295](#)
- rename iscsi target [221](#)
- rename logical device [91](#)
- rename SRP [158](#)
- rename storage group [117](#)
- rename storage groups [286](#)
- rename thin pools [74](#)
- replacing hba [295](#)
- report compression conversion status [155](#)
- report compression efficiency [154](#)
- report data compressibility [152](#)
- report flag details [210](#)
- report iscsi information [227](#)
- report non-native wwn [341](#), [342](#)
- report PE and VVol devices [245](#)
- report port information for iscsi target [231](#)
- report RP device types [264](#)
- report storage container for VVols [243](#)
- reservations
 - committing changes to [192](#)
 - device [191](#)
 - releasing [193](#)
- restore access control backup file [70](#)
- restore session with donor update [437](#)
- restore the authorization database [46](#)
- restoring session [437](#)
- restrictions
 - for deleting devices [190](#)
- Revert UnknwGrp to default [67](#)
- Role [37](#)
- RP device restrictions [248](#)

S

- scan for devices [31](#)
- SCSI device list [306](#)
- SCSI protocols [206](#)
- SCSI-level data [305](#)
- SDR
 - mapping [194–196](#)
- Service Level demand [166](#)
- session rollback [278](#)
- set array attributes [176](#)
- set CHAP credentials [296](#)
- set compression [150](#), [151](#)
- set device attributes [188](#)
- set HBA flags [293](#)
- set initiator group override flag [292](#)
- set online/offline state iscsi [227](#)
- set ORS session pace [428](#)
- set Service Level name [177](#)
- set SRP for storage group [102](#)
- set storage group SL [102](#)
- Set user authorization [42](#)
- show access control [65](#)
- show clone state flags [346](#)
- show device details [345](#)
- show device group details [88](#)
- show disk geometry details [346](#)

- show edisk information [271](#)
- show iscsi targets [231](#), [232](#)
- show Service Level [160](#)
- show SRP [164](#)
- show storage container for vVols [243](#)
- show thin pool details [351](#)
- show thin pool rebalancing [335](#)
- SMC Unisphere permissions [69](#)
- Solutions Enabler overview [24](#)
- SRDF device permissions [69](#)
- SRDF group attributes
 - setting [200](#)
- SRDF/A
 - assigning session priority [200](#)
- SRP reporting [165](#)
- start drain on edisk [273](#)
- stop background operations on thin devices [73](#)
- stop edisk drain [274](#)
- stop i/o activity [174](#)
- storage group demand [167](#)
- storage group details
 - grouping [113](#)
- storage groups [104](#)
- storgnsd, see also <Default para font> GNS, daemon options [141](#)
- support information [20](#)
- Symaccess permissions [69](#)
- SYMCLI commands
 - symcg [137](#)
 - symdg [123](#), [137](#)
 - symevent [369](#)
 - symrcopy query [431](#)
- symcli device reference [25](#)
- symcli help options [24](#)
- symconfigure restrictions [226](#)
- symsan command
 - ORS [419](#)

T

- Take front-end directors offline [322](#)
- Take RA directors offline [322](#)
- terminate and restart ORS session [425](#)
- terminate recoverpoint session [435](#)
- terminate session [434](#)
- thin device i/o activity [188](#)
- thin devices
 - converting [187](#)
- time-saving [27](#)
- timefinder permissions [69](#)
- truncated [27](#)

U

- update configuration data [35](#)
- unmap devices from front-end ports [197](#)
- unprotect journal devices [257](#)
- unprotect storage group devices [257](#)
- unsupported PE operations [247](#)
- Updating the host [199](#)
- user authorization [37](#)
- User authorization [42](#)
- user authorization command file [45](#)
- user ID formats [38](#)
- username groupname formats [39](#)

V

- verify array configuration [174](#)
- verify admin authority [55](#), [58](#), [60](#)
- verify auto-provisioning database [297](#)
- verify device usage [353](#)
- verify edisk state [273](#)
- verify ORS session [432](#)
- verify pool and device states [74](#)
- verify session checks [172](#)
- verify thin and data device states [352](#)
- view access control [65](#)
- view all thin device pools [75](#)
- view application registrations [314](#)
- view array metricst [177](#)
- view device service states [344](#)
- view external spindle information [365](#)
- View group details [300](#), [301](#)
- view hba alias name [302](#)
- view thin device allocations [79](#)
- view thin device details [78](#)
- view thin device pools details [76](#)
- view thin devices [77](#)
- virtual disk mapping [32](#)
- virtual provisioning overview [71](#)
- VLOGIX permission [55](#)
- vvols overview [237](#)
- vWitness management [202](#)

X

- XML [375](#)
- XML mode [375](#)
- XML output [376](#)
- XML output overview [375](#)