EMC ScaleIO

Version 2.0.x

User Guide

P/N 302-003-286

REV 04

EMC²

Copyright © 2016-2017 Dell Inc. or its subsidiaries. All rights reserved.

Published November 2017

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

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

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

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

Preface

As part of an effort to improve its product lines, 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 EMC technical support professional if a product does not function properly or does not function as described in this document.

Note

This document was accurate at publication time. Go to EMC Online Support (https://support.emc.com) to ensure that you are using the latest version of this document.

Related documentation

The release notes for your version includes the latest information for your product.

The following EMC publication sets provide information about your ScaleIO or ScaleIO Ready Node product:

- ScaleIO software (downloadable as ScaleIO Software <version> Documentation set)
- ScaleIO Ready Node with AMS (downloadable as ScaleIO Ready Node with AMS Documentation set)
- ScaleIO Ready Node no AMS (downloadable as ScaleIO Ready Node no AMS Documentation set)
- VxRack Node 100 Series (downloadable as VxRack Node 100 Series Documentation set)

You can download the release notes, the document sets, and other related documentation from EMC Online Support.

Typographical conventions

EMC uses the following type style conventions in this document:

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)
Italic	Used for full titles of publications referenced in text
`Monospace`	Used for: • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
Monospace italic	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

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

Product information

For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at https://support.emc.com.

Technical support

Go to EMC Online Support and click Service Center. You will see several options for contacting EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to techpubcomments@emc.com.

CONTENTS

Preface 3

Figures 17

Tables 19

Part 1 Introduction 23

Chapter 1 Introduction to EMC ScaleIO 25
What is ScaleIO? 26
System requirements 26
ScaleIO cluster components 27
Physical server requirements 27
Supported operating systems 28
GUI server requirements 30
ScaleIO Gateway server requirements 30
Other requirements 31
What's new in this version? 31
Product limits 32

Chapter 2 Architecture 33
ScaleIO Architecture Overview 34
System 34
Hardware 34
Software 34
The MDM cluster 35
Storage definitions 37
Protection Domains 37
Storage Pools 37
Fault Sets 38
Naming 39
Protection and load balancing 40
Rebuild 40
Rebuild throttling 40
Rebalance 41
Rebalance throttling 41
Checksum protection 42
ESX vStorage APIs for Array Integration (VAAI) 43
Caching 44
Networking 46
VMware limitation: 50
Virtual IP Address 51
S.M.A.R.T. hardware monitoring 51
List of approved RAID controllers 52
Snapshots 53
V-Trees 54
Other functions 55

Implementing ScaleIO....58
Physical layer.... 58
SAN virtualization layer.... 60
Implementing ScaleIO over a virtual system.... 62
VMware implementation....62
Xen implementation....64
Maintenance....64
Maintaining the physical layer....65
Instant maintenance mode....65
Maintaining the virtualization layer.... 66
Management tools....66
Configuring direct attached storage (DAS).... 67

Part 2 Getting Started 69

Chapter 3 Licensing 71
Overview.... 72
Activating entitlements and installing a license file.... 73
Activating an entitlement and downloading the license file.... 74
Installing the license....77
License file example....77
Error messages....78

Chapter 4 User Management 79
Overview....80
User Roles....80
Logging in.... 81
login....82
logout.... 83
Setting the User Authentication Method....83
set_user_authentication_method....83
Adding and modifying local users....84
add_user....84
delete_user.... 85
modify_user.... 85
query_users....86
query_user.... 86
reset_password.... 87
set_password.... 87
disable_admin....88
Resetting the admin user password.... 88

Chapter 5 Creating and Mapping Volumes 91
Creating and mapping volumes overview....92
Creating volumes....92
Adding volumes.... 92
Mapping a volume to an SDC....94
Mapping volumes....94

Part 3 Managing and Monitoring 97

Chapter 6 **Command Quick Reference** **99**
Device activities and monitoring..... 100
MDM activities and monitoring..... 101
Protection Domain activities and monitoring..... 103
SDC activities and monitoring..... 104
SDS activities and monitoring..... 106
Storage Pool activities and monitoring..... 108
User management activities and monitoring..... 110
Component authentication activities and monitoring..... 111
Volume activities and monitoring..... 112
Cache activities and monitoring..... 113
Other activities and monitoring..... 115
CLI basics..... 116
Using SCLI in non-secure mode..... 117
Syntax..... 118
Getting help with the CLI..... 119

Chapter 7 **Managing System Objects** **121**
Extend an existing ScaleIO system..... 122
Managing the MDM cluster..... 122
Replacing, or updating an IP address on a member of the MDM cluster..... 122
Configure virtual IP addresses..... 129
Managing SDC access to the MDM..... 129
Add another IP address subnet to an MDM cluster..... 129
Managing the SDSs and cache..... 131
Modifying an SDS port during IO..... 131
Managing Read Flash cache..... 133
Managing RAM read cache..... 134
SDC operations..... 134
Updating the SDC driver with IP changes..... 134
Detecting new volumes..... 135
Query volumes using drv_cfg..... 136
Query tgt objects using drv_cfg..... 136
Query GUID using drv_cfg..... 137
Query MDMs using drv_cfg..... 137
Loading a configuration file using drv_cfg..... 138
Adding an MDM using drv_cfg..... 139
Modifying an MDM IP address using drv_cfg..... 140
Permanent Device Loss state..... 141
Managing ESX servers..... 142
Modifying parameters on ESX servers..... 142
Checking the SDC state on ESX servers..... 144

Chapter 8 **Security Management** **147**
Setting up SSH authentication on the ScaleIO Gateway..... 148
Configuring SSL component authentication..... 148
Internal component authentication..... 148
External component authentication..... 148
Workflow for self-signed security certificates..... 149
Workflow for externally signed security certificates..... 149
Using Keytool to add certificates to external components..... 150
Managing SDC access to the MDM..... 152
Approved encryption methods..... 153

Login banner overview..... 153
Setting up a login banner using the CLI..... 154
Enabling or disabling preemptive acceptance of the login banner.... 154
Activating preemptive acceptance of the login banner.....155

Chapter 9 Opening the GUI and Logging In 157
Log in to the ScaleIO GUI.....158
Connection and disconnection information..... 158

Chapter 10 GUI Features 159
GUI overview..... 160
GUI conventions.....160
Alerts indicators.....161
Color codes.....161
Dashboard view.....162
Navigation tree.....163
Preferences menu.....164
Dimmer.....164
Widget..... 164
Dashboard tiles..... 165
Management (MDM) cluster status..... 168
Frontend views.....170
Volumes..... 171
SDCs..... 172
Snapshots..... 172
Adding Frontend objects to the filter.....172
Backend view..... 173
Command menu (3).....175
Filter (1)..... 175
Table..... 175
Alerts view..... 179
Property Sheets.....180

Chapter 11 Monitoring the System using the GUI 185
Viewing object properties in Backend and Frontend views..... 186
Viewing licensing information..... 186
Verifying your connection to the Management cluster..... 186

Chapter 12 Configuring the System using the GUI 187
Configuring capacity..... 188
Adding SDSs and storage devices..... 188
Removing SDSs and devices..... 191
Adding, removing, and activating and inactivating capacity.....192
Activating devices..... 194
Clearing device errors..... 194
Setting device capacity limits.....195
Configuring cache.....195
Setting Read Flash Cache policy at Storage Pool level.....195
Setting Read Flash Cache policy at SDS level..... 196
Adding Read Flash Cache devices..... 197
Changing RAM Read Cache volume settings.....197
Removing Read Flash Cache devices..... 198

Configuring RAM Read Cache (advanced, Backend).....198
Configuring volumes, volume trees, SDCs, and snapshots.....200
Adding volumes..... 200
Mapping and unmapping volumes..... 202
Removing volumes.....204
Creating volume snapshots.....205
Removing snapshots.....206
Removing snapshots from a consistency group..... 207
Increasing a volume's size.....208
Setting volume bandwidth and IOPS limits..... 209
Entering and exiting SDS Instant Maintenance Mode..... 210
Configuring and viewing Oscillating Failure counters..... 211
Configuring Oscillating Failure counter parameters.....211
Resetting Oscillating Failure counters..... 214
Viewing Oscillating Failure counters..... 215
Applying Performance Profiles to system components..... 216
Configuring I/O priorities and bandwidth use (advanced)..... 217
Application IOPS and bandwidth (advanced).....217
System IOPS and bandwidth (advanced)..... 217
Enabling and disabling Rebuild\Rebalance (advanced).....219
Using the background device scanner.....219
Enabling and disabling the background device scanner.....220
Resetting the background device scanner counters..... 221
Modifying Checksum protection mode.....221
Renaming objects..... 222
Approving pending security certificates.....223
Customizing system preferences.....223

Chapter 13 Using the VMware Plug-in 227
VMware Plug-in overview.....228
Configuring components..... 229
Configuring components—basic..... 231
Configuring components—advanced.....232
Viewing components..... 238

Part 4 Reference 241

Chapter 14 Troubleshooting 243
Troubleshooting ScaleIO.....244
High latency encountered when S.M.A.R.T. hardware monitoring feature is enabled.....244
Adding a cache device to RFcache Storage Pool.....244
Add SDS device with vSphere plug-in fails in DirectPath environment..... 247
After SDC installation, the ScaleIO SVM does not start automatically.....247
Application server does not see a ScaleIO volume..... 248
Cannot log in to the Installation Manager after upgrade..... 248
Certificate error when installing SDC on Windows servers..... 249
VMware deployment failures 249
Enabling acceleration in a new node..... 250
S.M.A.R.T. hardware alerts are not displayed.....253
Gateway server recovery.....253
Recovering Gateway on a new server during upgrade..... 253
Recovering IM with LIA trusted IP feature enabled.....255

SNMP configuration recovery post gateway crash during upgrade... 256
Installation Manager returns an error....257
Installation with the Installation Manager fails.... 257
LIA upgrade fails in Windows 2008.... 257
Error when unmapping volume from ESXi server....258
Older SDCs cannot communicate with newer SDSs after upgrade to v2.x 258
Removing RFcache from Windows servers....259
Removing RFcache leftovers from the Windows OS registry.... 259
Replacing a faulty caching device in ScaleIO.... 259
Replacing a faulty accelerated storage device in ScaleIO....261
ScaleIO CLI or GUI cannot connect to an MDM....262
ScaleIO Gateway fails to run.... 263
SCLI add_sds command fails due to communication error or MDM going offline.... 263
SVM manual memory allocation....263
Viewing the status of volumes by drv_cfg --rescan command....265
Virtual IP feature is not functional.... 266
Enabling acceleration in a new node.... 267
The VMware plug-in responds slowly.... 269
Solving ScaleIO performance issues.... 270
Mismatch in IO counters....270
Deploying SVM on a node with a management IP address from a different subnet than the node with the SVM template.... 271
Speeding up rebuild and rebalance processes....271
RFcache is not installed on an upgraded Windows server....271
SSD devices are not recognized in the operating system....272
ESX competing thread setting resets on server reboot.... 273

Chapter 15 Frequently Asked Questions 275
Install the ScaleIO GUI....276
Associating ScaleIO volumes with physical disks.... 276
Volume information - Linux.... 276
Volume information - Windows.... 277
Port usage and changing default ports.... 278
Adding an external SDC to an existing ScaleIO system....279
Installing SDC on an ESX server and connecting it to ScaleIO..... 279
Installing SDC on a Linux server and connecting it to ScaleIO..... 280
Installing SDC on a Windows server and connecting it to ScaleIO.... 281
Changing the LIA configuration file.... 281
Cleaning the ScaleIO VMware environment and performing a clean install.... 282
Configuring ScaleIO devices in Linux LVM....283
Configuring session timeout parameters....284
Fixing keytool errors.... 284
Error during rpm installation command.... 285
Error during rpm upgrade command.... 285
Installing Java on SUSE 12 servers.... 285
Mounting ScaleIO....286
The ScaleIO Gateway web server isn't responding.... 287
The ScaleIO Gateway (REST service, Installation Manager) may be disabled:.... 287

The ScaleIO Gateway web server isn't responsive and the following error appears in the catalina log file:.....288
Upgrading the Gateway when a custom certificate is used.....289
Uploading a new OVA.....289
Using the same data network for different NICs.....289
What to do when the default self-signed certificate expires.....289
Add another IP address subnet to an MDM cluster.....290

Chapter 16 System events 293
System events overview.....294
Event format.....294
Viewing events.....295
Viewing events locally.....295
Event list.....298
Authentication.....298
CLI commands.....299
License and installation.....300
MDM.....302
SDC.....304
SDS.....305
Rebuild.....308

Chapter 17 Return messages 311
Return messages overview.....312
Message list.....312
SDS device name changed.....312
Success.....312
No resources.....313
Timeout.....313
Already exists.....313
Communication error.....314
Wrong reconfiguration mode.....314
SDS add device “open” error.....314
SDS add device size problem.....314
SDS add device IO error.....315
SDS actively attach to this MDM.....315
SDS already attached to this MDM.....315
SDS active attach to other MDM.....316
SDS already attached to other MDM.....316
SDS not found.....317
Volume not found.....317
SDC not found.....317
Volume already mapped to SDC.....318
Volume already mapped to all SDCs.....318
Volume not mapped to SDC.....318
Volume mapped to all SDCs.....319
Too many mappings.....319
Invalid volume ID.....319
Invalid SDS ID.....320
Invalid SDC ID.....320
Host not found.....320
Duplicate host name.....321
Address and port in use.....321
Too many SDSs in Protection Domain.....321

Too many SDSs....321
Too many volumes in Protection Domain....322
Too many volumes....322
SDS name in use....322
Volume name in use....323
Duplicate SDS device name....323
Duplicate SDC IP address....323
SDS not under removal....324
SDS being removed....324
Volume mapped....324
Use MDM IP address....325
Command only in cluster....325
Command only in single mode....326
Command error when cluster degraded....326
Too many devices....326
SDS device not found....327
License too long....327
License error....327
Unknown license version....328
License not set....328
Configuration exceeds new license....328
License capacity exceeded....329
License expired....329
Installation ID mismatch....329
Customer ID mismatch....330
Unsupported license version....330
Volume decrease not supported....330
SDS device being removed....331
Cannot remove last SDS device....331
Device not under removal....331
SDS already adding device....332
Invalid device ID....332
Too many Protection Domains....332
Protection Domain name in use....333
Protection Domain not found....333
Protection Domain has SDSs....333
Not enough SDSs in Storage Pool....334
Remove in progress....334
No network test results....334
Network test in progress....335
SDS device already active....335
SDS device wrong test mode....335
SDS device wrong test state....336
SDS device test in progress....336
Volume allocation “busy” error....336
Capacity highly unbalanced....337
Volume creation failed....337
Invalid password....338
No permissions....338
Too many Storage Pools in Protection Domain....338
Storage Pool already exists....339
Storage Pool not found....339
Storage Pool has devices....339
Storage Pool name already exists....340
Invalid Storage Pool ID....340
Protection Domain has Storage Pools....340

SDC has mappings..... 341
Too many volumes in Vtree..... 341
Too many snapshots in Vtree..... 341
No volume to delete.....342
Too many devices in Storage Pool..... 342
Volume not found for consistency group..... 342
Too many IP addresses..... 343
Cannot remove last SDS IP address.....343
SDS IP address not found.....343
Volume size too large.....344
Capacity too low for snapshot..... 344
IP address belongs to other SDS..... 344
Invalid port.....345
LIA package not installed..... 345
LIA multiple packages installed..... 345
LIA command failed..... 346
LIA invalid password..... 346
LIA operation not permitted..... 347
LIA file is too large..... 347
LIA wrong ID..... 347
LIA disk does not exist..... 348
SDS IP address not found.....348
Bad threshold value..... 348
Storage Pool has Volumes..... 349
Protection Domain and Storage Pool mismatch.....349
SDS IP address already exists.....349
IP role already set.....350
Cluster ID mismatch..... 350
Device already exists..... 350
Volume size is illegal..... 351
SDC must be disconnected..... 351
Invalid percentage..... 352
Wrong receive group..... 352
Mismatched SW ID..... 352
Invalid license issuer..... 353
Number of parallel messages is too low..... 353
Number of parallel messages is too high..... 353
Network test size is too high..... 354
Network test length too long..... 354
Network test must limit either time or size..... 354
Not enough SDSs in Protection Domain for test.....355
Not allowed in Storage Pool with devices..... 355
License capacity mismatch.....355
Exceeds system capacity limitations.....356
SDC mapped to too many volumes..... 356
Not enough devices in Storage Pool for removal..... 357
Invalid network limits..... 357
A device is currently being removed..... 357
A device is already active.....358
RAM Cache wrong state.....358
More than one mapping not allowed..... 359
At least one SDS failed..... 359
A device is in incorrect test state..... 359
A device test is in progress..... 360
Enterprise features not enabled.....360
Volume already mapped to an SDC.....361

SDS already exists and is being removed....361
Invalid spare percentage....362
Invalid background device scanner bandwidth limit....362
Zero padding with background device scanner data comparison mode....362
IP address already assigned to another SDS....363
Capacity too low for thin volumes....363
Too many Storage Pools in the system....363
Remote access to the MDM is blocked....364
Cannot add a partitioned device....364
Cannot add a mounted device....365
Device header mismatch....365

Chapter 18 ScaleIO on Xen 367
Overview of ScaleIO on Xen....368
Adding a volume in XenServer environment....368
Removing a ScaleIO volume from Xen....369
Modifying the size of a ScaleIO volume....370
Xen v6.5 High Availability....371

Chapter 19 Configuring ScaleIO in OpenStack Environments 373
Overview....374
Volume limitations....374
Supported operations....374
QoS functionality....374
Thin provisioning....375
Multiple Protection Domains and Storage Pools....375
Static configuration mode....375
Dynamic mode....375
Pool-aware scheduler support....377

Chapter 20 REST API Reference 379
Before you begin....380
Configuring connection properties in the ScaleIO Gateway....380
Logs....384
CLI/REST API command reference....384
Software type resources....390
General information....390
System....404
Protection Domain....438
SDS....452
Storage Pool....471
Device....485
Volume....494
VTree....499
SDC....500
User Objects....504
Fault Set....506
RFcache Device....510
Alert Objects....512
ESRS Objects....515
Responses....516
Output Content Format....516
Return Codes....516

Changes that might require adjustments of old code..... 517
Deleted scsiInitiator commands.....518
Other deleted commands..... 518
Cluster commands..... 518
Deprecated alerts.....519
Using the IM REST API..... 519
General information.....519
How to use IM REST API URIs.....523
Deploying a system.....524
Installation Manager commands..... 525

Chapter 21 SNMP trap support 533
General.....534
Supported alerts and event numbering conventions..... 534
CLASS/TYPE:..... 534
ScaleIO Alerts in SNMP, GUI, REST, and ESRS..... 535
Configure SNMP properties after deployment..... 544
Create a Lockbox.....544
Configuring SNMP after deployment..... 545
ScaleIO.mib file..... 546

Glossary 549

FIGURES

1 5-node MDM cluster.... 36
2 Protection Domains and Storage Pools....37
3 Protection Domains, Storage Pools, and Fault Sets....39
4 ScaleIO system deployed on a single network.... 49
5 ScaleIO system deployed on separate networks....50
6 Snapshot operations....54
7 V-Tree diagram.... 55
8 Physical layout example—3-node cluster....60
9 ScaleIO implementation on ESX.... 62
10 ScaleIO Xen virtual machine architecture.... 64
11 Licensing LAC email....73
12 License file example....77
13 Add Volume window.... 94
14 Map Volumes window after mapping is complete.... 95
15 Alerts indicators....161
16 Dashboard controls.... 163
17 Widget, showing Return to Full-Screen button.... 165
18 Dashboard tiles.... 166
19 Frontend > Volumes view....171
20 Frontend > SDCs view....172
21 Frontend > filtered Snapshots view 172
22 Backend view.... 174
23 Backend filter....175
24 Alerts view.... 179
25 Multiple floating Property Sheets....182
26 Example of a Property Sheet for an SDS....183
27 Add SDS window....190
28 Add Device window showing command validation....191
29 Set Volume Use RAM Read Cache window.... 198
30 RAM Read Cache configuration at Storage Pool level....199
31 RAM Read Cache configuration at SDS level....200
32 Add Volume window.... 202
33 Map Volumes window after mapping is complete.... 203
34 Unmap Volumes window....203
35 Remove Volumes window.... 204
36 Snapshot Volume window....206
37 Removal of a volume and associated snapshots....207
38 Removal of snapshots of a specific volume....207
39 Remove Consistency Groups window....208
40 Increase Volumes' Size window....209
41 Set Volume Limits window.... 210
42 Configure Oscillating Failure counters—System....213
43 Configure Oscillating Failure counters—Protection Domain or Storage Pool.... 213
44 Reset Oscillating Failure counters—System.... 214
45 Configure Oscillating Failure counters—Protection Domain or Storage Pool.... 215
46 Set Performance Profile window....217
47 Set Network Throttling window.... 219
48 Background device scanner configuration....221
49 Reset Background Device Scanner Counters window.... 221
50 Configure Use Checksum window.... 222
51 Preferences window....224
52 Configure virtual IPs dialog box.... 238
53 Example of IM REST API URI when JSESSION ID is configured as a header.... 524

TABLES

1 Server physical requirements....27
2 Supported operating systems - ScaleIO components.... 29
3 Product limits.... 32
4 MDM cluster modes.... 36
5 Caching modes.... 44
6 Caching support matrix.... 46
7 IP address configurations in ScaleIO (based on CSV file).... 47
8 eLicensing terminology(continued)....73
9 Licensing error messages(continued).... 78
10 Local and LDAP user roles and permissions.... 80
11 Device activities and monitoring....100
12 MDM activities and monitoring.... 101
13 Protection Domain activities and monitoring.... 103
14 SDC activities and monitoring.... 104
15 SDS activities and monitoring.... 106
16 Storage Pool activities and monitoring....108
17 User management activities and monitoring.... 110
18 Component authentication activities and monitoring....111
19 Volume activities and monitoring.... 112
20 Cache activities and monitoring.... 113
21 Other activities and monitoring....115
22 GUI task overview....160
23 GUI color codes.... 161
24 Alert symbols and color codes....162
25 Dashboard tiles.... 166
26 Management node icons with normal operational status (green).... 169
27 Management node status indications and color codes....170
28 Table view options....176
29 Preferences....224
30 Plug-in activity matrix....229
31 Default ports....278
32 ScaleIO CLI commands....384
33 Response.... 391
34 System Objects.... 404
35 System Relationships.... 406
36 MDM Cluster Object.... 406
37 MdmClusterMember....407
38 MdmPerfParams.... 407
39 OscillatingCounterParameters....408
40 OscillatingCounterWindowParameters.... 408
41 System Statistics Object....408
42 System Bandwidth Counters 414
43 System Operations....415
44 Protection Domain Object.... 438
45 Protection Domain Parent.... 439
46 Protection Domain Relationships....439
47 Protection Domain Statistics Object....439
48 Protection Domain Operations....445
49 SdsConnectionState.... 451
50 SdsPeerConnectionState....451
51 SingleConnectionState....452
52 DeviceTestResults....452
53 DeviceTestResult.... 452

54 SDS Object....452
55 SDS Parent....454
56 SDS Relationships....454
57 Sdslp....455
58 SdsRfcacheDeviceInfo....455
59 CertificateInfo....455
60 SdsPerfParams....455
61 SDS Statistics Object....456
62 SDS Operations....462
63 NetBwc....470
64 MosBwcQueryWindowSize....471
65 MosBwcQueryWindowLatency....471
66 Storage Pool Object....471
67 Storage Pool Parent....472
68 Storage Pool Relationships....473
69 Storage Pool Statistics Object....473
70 Storage Pool Operations....477
71 Device Object....485
72 Device Parent....486
73 Device Relationships....486
74 OscillatingCounter....486
75 OscillatingCounterWindow....486
76 Device Statistics Object....487
77 Device Operations....490
78 Volume Object....494
79 Volume Parent....495
80 Volume Relationships....495
81 Volume Statistics Object....495
82 SdcMappingInfo....495
83 Volume Operations....495
84 VTree Object....499
85 VTree Parent....499
86 VTree Relationships....500
87 VTree Statistics Object....500
88 SDC Object....500
89 SDC Parent....501
90 SDC Relationships....501
91 SdcPerfParams....501
92 SDC Statistics Object....502
93 SDC Operations....502
94 User Object....504
95 User Parent....505
96 User Operations....505
97 Fault Set Object....506
98 Fault Set Parent....506
99 Fault Set Relationships....506
100 Fault Set Statistics Object....506
101 Fault Set Operations....508
102 RFcache Device Objects....510
103 RFcache Device Parent....510
104 rfcacheDevice Statistics Object....510
105 rfcacheDevice Operations....511
106 Alert Object....512
107 Alert Operations....515
108 ESRS Operations....515
109 Return codes....516

110 Operations.....525
111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS..... 535

PART 1

Introduction

The section describes an overview of the benefits and system requirements of ScaleIO. Chapters include:

Chapter 1, "Introduction to EMC ScaleIO"

Chapter 2, "Architecture"

CHAPTER 1

Introduction to EMC ScaleIO

This section introduces EMC ScaleIO. Topics include:

- What is ScaleIO? 26
- System requirements 26
- What's new in this version? 31
- Product limits 32

What is ScaleIO?

ScaleIO

ScaleIO is a software-only solution that uses existing servers' local disks and LAN to create a virtual SAN that has all the benefits of external storage—but at a fraction of cost and complexity. ScaleIO utilizes the existing local storage devices and turns them into shared block storage. For many workloads, ScaleIO storage is comparable to, or better than external shared block storage.

The lightweight ScaleIO software components are installed on the application servers and communicate via a standard LAN to handle the application I/O requests sent to ScaleIO block volumes. An extremely efficient decentralized block I/O flow, combined with a distributed, sliced volume layout, results in a massively parallel I/O system that can scale up to thousands of nodes.

ScaleIO is designed and implemented with enterprise-grade resilience. Furthermore, the software features an efficient distributed self-healing process that overcomes media and server failures, without requiring administrator involvement.

Dynamic and elastic, ScaleIO enables administrators to add or remove servers and capacity on-the-fly. The software immediately responds to the changes, rebalancing the storage distribution and achieving a layout that optimally suits the new configuration.

Because ScaleIO is hardware agnostic, the software works efficiently with various types of disks, including: magnetic (HDD) and solid-state disks (SSD), flash PCI Express (PCIe) cards, networks, and hosts.

ScaleIO can easily be installed in an existing infrastructure as well as in green field configurations.

ScaleIO Ready Node

ScaleIO Ready Node is the combination of ScaleIO software-defined block storage and Dell PowerEdge® servers, optimized to run ScaleIO, enabling customers to quickly deploy a fully architected, software-defined, scale out server SAN.

Any hypervisor can run on ScaleIO Ready Node servers as an application consuming ScaleIO volumes.

In an AMS-based solution, a limited number of operating systems are currently supported. For more information, see the supported OS tables.

For VMware, the converged solution is managed by the AMS (Automated Management Services), which enables a simple or customized deployment process from bare metal, no IP state, to a fully-configured system: IP address assignment, ScaleIO deployment and configuration, and vCenter configuration.

System requirements

This section lists the requirements for system components.

This section is specific to ScaleIO software deployments.

For ScaleIO Ready Node or VxRack Node 100 Series systems, refer to your product's Hardware Configuration and Operating System Installation Guide.

ScaleIO cluster components

List of required ScaleIO servers.

- ScaleIO component servers:
  - 3-node cluster
    - One Master MDM
    - One Slave MDM
    - One Tie Breaker
    - Minimum of three SDSs (on the same servers as the above components, or on three different servers)
    - SDCs, up to the maximum allowed (on the same servers as the above components, or on different servers)
  - 5-node cluster
    - One Master MDM
    - Two Slave MDMs
    - Two Tie Breakers
    - Minimum of three SDSs (on the same servers as the above components, or on three different servers)
    - SDCs, up to the maximum allowed (on the same servers as the above components, or on different servers)
- ScaleIO Gateway server on a separate server, or together with an MDM or SDS. Do not install the Gateway on an SDC server or on an SDS on which RFcache will be enabled.
- ScaleIO Gateway server on a separate server, outside of the ScaleIO system.

Physical server requirements

Table 1 Server physical requirements

Component	Requirement
Processor	One of the following: • Intel or AMD x86 64-bit (recommended) • Intel or AMD x86 32-bit (for Xen only)
Physical memory	ScaleIO component requirements: • 500 MB RAM for the Meta Data Manager (MDM) • 500 MB RAM for each ScaleIO Data Server (SDS) • 50 MB RAM for each ScaleIO Data Client (SDC) DAS Cache memory requirements (ScaleIO Ready Node, non-XenServers only). Add to every SVM/node that will be using DAS Cache: • 1U1N servers—500 MB RAM • 2U1N servers—1 GB RAM

Table 1 Server physical requirements (continued)

Component	Requirement
	To calculate SVM memory allocation, use the formulas provided in the *ScaleIO Deployment Guide*.
Disk space	• 1 GB for each physical node or Xen hypervisor • 10 GB for VMware topologies

Connectivity	One of the following:
	• 1 GbE or 10 GbE (recommended) network
	• IP-over-InfiniBand network
	Dual-port network interface cards (recommended)
	Ensure the following:
	• There is network connectivity between all components.
	• Network bandwidth and latency between all nodes is acceptable, according to application demands.
	• Ethernet switch supports the bandwidth between network nodes.
	• MTU settings are consistent across all servers and switches.
	• The following TCP ports are not used by any other application, and are open in the local firewall of the server:
	  ▪ MDM: 6611 and 9011
	  ▪ Tie Breaker: 9011
	  ▪ SDS: 7072. Multiple SDS (not supported on Windows): 7073-7076
	  ▪ Light Installation Agent (LIA): 9099
	  ▪ SDBG ports (used by ScaleIO internal debugging tools to extract live information from the system): MDM 25620, SDS 25640. Multiple SDS (not supported on Windows): 25641-25644 (not 25640).
	• The following UDP port is open in the local firewall of the server:
	  ▪ SNMP traps: 162
	Note
	You can change the default ports. For more information, see “Changing default ports” in the user documentation.

Supported operating systems

The list of operating systems supported by this version of ScaleIO.

For the most updated list, see the EMC Simple Support Matrix (ESSM) at https://elabnavigator.emc.com/eln/elnhome.

Table 2 Supported operating systems - ScaleIO components

Operating system	Requirement
Linux	Supported versions: • CentOS 6.x-7.x, Oracle Linux 6.5/7.x • Red Hat 6.x-7.x • SUSE 11.3, 11.4, 12, 12.1, 12.2 • Ubuntu 14.04, Ubuntu 16.04 **Note** Before deploying on Ubuntu servers, you must prepare the environment, as described in the *EMC ScaleIO Deployment Guide*. Packages required for all components, all Linux flavors: • numactl • libaio Additional packages required for MDM components: • bash-completion (for SCLI completion) • Latest version of Python 2.X When installing the MDM component on Linux CentOS 6 or RHEL 6 hosts (for software-only systems), set the shared memory parameter in the /etc/sysctl.conf file to at least the following value: kernel.shmmax=209715200. To use this value, type the `sysctl -p` command. To use the secure authentication mode, ensure that OpenSSL 64-bit v1.0.1 or later (v1.1, however, is not supported) is installed on all servers in the system. To use the secure authentication mode on SUSE 11.3/11.4 servers, ensure that the OpenSSL on the server is v1.0.1 or later (v1.1, however, is not supported), or install these packages (from the ISO in the Complete VMware SW download container) on the server: • `libopenssl1_0_0-1.0.1g-0.40.1.x86_64.rpm` • `openssl1-1.0.1g-0.40.1.x86_64.rpm` To use LDAP, ensure that OpenLDAP 2.4 is installed on all servers.
Windows	Supported versions: • 2008 R2, 2012, 2012 R2, or 2016 (in v2.0.1.1 and later). Server Core editions are not supported. • For VxRack Node 100 Series, only 2012 R2 is supported. • On all MDM servers, install the EMC-provided `PythonModulesInstall.exe` on all MDM nodes. The file is supplied on the ISO, or download from the EMC Online Support site (search for ScaleIO Python Installation Modules) on https://support.emc.com. • To install SDC or RFcache on 2008 R2, ensure that Microsoft Security Update KB3033929 is installed. To use the secure authentication mode, ensure that these are installed on all servers in the system: • OpenSSL 64-bit v1.0.1 or later (v1.1, however, is not supported)

Table 2 Supported operating systems - ScaleIO components (continued)

Operating system	Requirement
	• Visual C++ redistributable 2010 package, 64-bit To use RFcache, ensure that Visual C++ redistributable 2010 package, 64-bit is installed on all servers in the MDM cluster and on all SDSs.
Hypervisors	• VMware ESXi OS: 5.5 U3, 6.0 U3, or 6.5, managed by vCenter 5.5 or 6.0 • Hyper-V • XenServer 6.5 or 7.0 **Note** OpenSSL 64-bit v1.0.1 is supported on XenServer 6.5 SP1 (or later) • Red Hat KVM

GUI server requirements

Component	Requirement
Supported operating systems	Windows 7, 2008 R2, 10, 2012 or 2012 R2, 2016 (in v2.0.1.1 and later). Server Core editions are not supported.
Other	• Java 1.8 64-bit or later • Screen resolution: 1366 x 768 minimum

ScaleIO Gateway server requirements

Component	Requirement

Supported operating systems	• Windows 2008 R2, 2012 R2, or 2016, including the Visual C++ redistributable 2010 package, 64-bit. Server Core editions are not supported.
	• Linux:
	  ▪ CentOS 6.x-7.x
	  ▪ Oracle Linux 6.5/7.x
	  ▪ Red Hat 6.x-7.x
	  ▪ SUSE 11.3, 12, 12.1, and 12.2
	  ▪ Ubuntu 14.04, Ubuntu 16.04
	Every server requires 2 cores and a minimum of 3 GB available RAM.

Connectivity	The following TCP ports are not used by any other application, and are open in the local firewall of the server: 80 and 443 (or 8080 and 8443). You can change the default ports. For more information, see “Changing default ports” in the user documentation.
Supported web browsers	• Internet Explorer 10, or later

Component	Requirement
	• Firefox, version 42, or later • Chrome, version 45, or later
Java requirements	• 1.8 or later, 64-bit
Other	• For a Windows Gateway, the Windows Management Instrumentation service must be enabled on the IM server and on all Windows ScaleIO nodes. • Do not install the Gateway on a server on which RFcache will be enabled or on which SDC will be installed. • The Gateway server must have connectivity to all the nodes that are being installed. If you are using separate networks for management and data, the server must be able to communicate with both networks.

Other requirements

ScaleIO requires that you use a minimum of three SDS servers, with a combined free capacity of at least 300 GB. These minimum values are true per system and per Storage Pool.

NOTICE

ScaleIO installation enables unlimited use of the product, in non-production environments . To obtain a license for production use, and to receive technical support, open a service ticket with Customer Support at https://support.emc.com.

For complete information on licensing, see the *ScaleIO User Guide*.

What's new in this version?

This version introduces the following new features:

- Devices attached to a RAID controller on an ESXi server can be managed with VMware vSphere VMDirectPath I/O (PCI passthrough, or DirectPath). This reduces the CPU impact of high-bandwidth throughput workloads by "bypassing" the hypervisor. For ScaleIO Ready Node servers, the ScaleIO vSphere plug-in performs all the functionality necessary for the server and device configuration. For other servers, enable DirectPath on the relevant PCI device of the SAS\Raid Controller in the vSphere client, then use the ScaleIO vSphere plug-in to add PCI devices to the ScaleIO SDS.
- Support has been added for VMware vSphere Update Manager (VUM). VUM is a software service for automating patch management. In v2.0.1.4, SDC for ESX will support VUM 6.0 and above.
  For additional information, refer to:
  - VUM details: http://searchvmware.techtarget.com/definition/VMware-vSphere-Update-Manager-VUM
  - VUM official documentation: https://www.vmware.com/support/pubs/vum_pubs.html

Product limits

The following table lists product capabilities:

Table 3 Product limits

Item	Limit
ScaleIO System raw capacity	300 GB—16 PB
Device size	100 GB—8 TB
Minimum Storage Pool capacity	300 GB
Volume size	8 GB—1 PB
Maximum number of volumes/snapshots in system	32,768[a]
Maximum number of volumes/snapshots in Protection Domain	32,768
Maximum number of volumes + snapshots in single VTree	32
Maximum capacity per SDS	96 TB
SDSs per system	1024
SDSs per Protection Domain	128[a]
Maximum devices (disks) per SDS server	64[b]
Maximum devices (disks) per Storage Pool	300[a]
Minimum devices (disks) per Storage Pool	3, on different SDSs
Maximum SDCs per system	1024 When using replication with RecoverPoint, the maximum number of SDCs is reduced by the number of RPAs in the system.[c]
Maximum volumes that can be mapped to a single SDC	8192
Maximum Protection Domains per system	256
Maximum Storage Pools	1024
Maximum Storage Pools per Protection Domain	64
Maximum Fault Sets per Protection Domain	64
Maximum IP addresses per server (MDM and SDS)	8
RAM Cache	128 MB—300 GB

a. If more are needed, contact EMC Support.
b. On VMware servers, the maximum devices per SDS is 59.
c. Replication support is version-specific. For information, see the ESSM.

CHAPTER 2

Architecture

This chapter describes the ScaleIO architecture. Topics include:

- ScaleIO Architecture Overview 34
- System 34
- The MDM cluster 35
- Storage definitions 37
- Protection and load balancing 40
- ESX vStorage APIs for Array Integration (VAAI) 43
- Caching 44
- Networking 46
- Virtual IP Address 51
- S.M.A.R.T. hardware monitoring 51
- Snapshots 53
- Other functions 55
- Implementing ScaleIO 58
- Implementing ScaleIO over a virtual system 62
- Maintenance 64
- Management tools 66
- Configuring direct attached storage (DAS) 67

ScaleIO Architecture Overview

This chapter describes the ScaleIO architecture overview.

ScaleIO is a software-only solution. ScaleIO components are lightweight, highly available software components, installed on new or existing servers alongside your production applications (hypervisors, databases, web applications, etc.). The system can be installed directly on the servers, or over a virtual server system (hypervisor or virtual machines).

System

The ScaleIO system is based on a hardware and a software component.

Hardware

In general, hardware can be the existing application servers used by the datacenter, or a new set of nodes (if, for example, you want to dedicate all nodes solely for the purpose of running the ScaleIO SAN storage system).

- *Nodes*
  Nodes, or servers, are the basic computer unit used to install and run the ScaleIO system. They can be the same servers used for the applications (server convergence), or a dedicated cluster. In any case, ScaleIO is hardware-agnostic, and therefore, aside from performance considerations, the type of server is inconsequential.
- *Storage Media*
  The storage media can be any storage media, in terms of the type (HDD, SSD, or PCIe flash cards) and anywhere (DAS, or external).

Software

The ScaleIO virtual SAN consists of the following software components:

- Meta Data Manager—MDM: Configures and monitors the ScaleIO system. The MDM can be configured in redundant Cluster Mode, with three members on three servers or five members on five servers, or in Single Mode on a single server.

  NOTICE

  It is not recommended to use Single Mode in production systems, except in temporary situations. The MDMs contains all the metadata required for system operation. Single Mode has no protection, and exposes the system to a single point of failure.

- ScaleIO Data Server—SDS: Manages the capacity of a single server and acts as a back-end for data access. The SDS is installed on all servers contributing storage devices to the ScaleIO system. These devices are accessed through the SDS.
- ScaleIO Data Client—SDC: A lightweight device driver that exposes ScaleIO volumes as block devices to the application that resides on the same server on which the SDC is installed.

Depending on the desired configuration (described later), the software components are installed on the server and give rise to a virtual SAN layer exposed to the applications residing on the servers.

The MDM cluster

The MDM serves as the monitoring and configuration agent of the ScaleIO system. The MDM is mainly used for management which consists of migration, rebuilds, and all system-related functions. No I/O run through the MDM.

To support high availability, three or more instances of MDM run on different servers. In a multi-MDM environment, one MDM is given the Master role, and the others act as Slave or Tie Breaker MDMs.

The MDM cluster comprises a combination of Master MDM, Slave MDMs, and Tie Breaker MDMs.

The following terms are relevant to the MDM, the building blocks of the MDM cluster:

- MDM

  Any server with the MDM package installed on it. An MDM can be given a Manager or a Tie Breaker (default) role, during installation. MDMs have a unique MDM ID, and can be given unique names.

  Before the MDM can be part of the cluster, it must be promoted to a Standby MDM.

- Standby MDM and Tie Breaker

  An MDM and a Tie Breaker can be added to a system as a standby. Once added, the standby MDM or Tie Breaker is attached, or locked, to that specific system.

  A standby MDM can be called on to assume the position of a Manager MDM or Tie Breaker MDM when it is promoted to be a cluster member.

- Manager MDM

  An MDM that can act as a Master or a Slave in the cluster. Manager MDMs have a unique system ID, and can be given unique names. A manager can be a standby or a member of the cluster.

  In ScaleIO documentation, “MDM” refers to a manager, unless specified otherwise.

- Tie Breaker MDM

  An MDM whose sole role is to help determine which MDM is the master. A Tie Breaker can be a standby or a member of the cluster. A Tie Breaker MDM is not a manager.

  In a 3-node cluster, there is one TB; in a 5-node cluster, there are two TBs. This ensures that there are always an odd number of MDMs in a cluster, which guarantees that there is always a majority in electing the master.

The following terms are relevant to the MDM cluster, specifically:

- Master MDM (used to be called Primary MDM)

  The MDM in the cluster that controls the SDSs and SDCs. The Master MDM contains and updates the MDM repository, the database that stores the SDS configuration, and how data is distributed between the SDSs in the system. This repository is constantly replicated to the Slave MDMs, so they can take over with no delay.

  Every MDM cluster has one Master MDM.

- Slave MDM (used to be called Secondary MDM)

  An MDM in the cluster that is ready to take over the Master MDM role if ever necessary.

In a 3-node cluster, there is one Slave MDM, thus allowing for a single point of failure. In a 5-node cluster, there are two Slave MDMs, thus allowing for two points of failure. This increased resiliency is a major benefit to enabling the 5-node cluster.

- Replica

  An MDM that contains a replica of the MDM repository. This includes the Master MDM and any Slave MDMs in the MDM cluster.

The following table describes the available cluster modes:

Table 4 MDM cluster modes

Cluster mode	Members	Description
3-node (default)	• Master MDM • Slave MDM • Tie Breaker	3-node cluster has two copies of the repository, thus can withstand one MDM cluster failure.
5-node	• Master MDM • Two Slave MDM • Two Tie Breaker	5-node cluster has three copies of the repository, thus can withstand two MDM cluster failure.
Single-node	• Master MDM	Single-node cluster has only one copy of the repository, thus it cannot withstand failure. It is not recommended to use Single Mode in production systems, except in temporary situations.

In addition to the cluster members, you can prepare standby Managers and Tie Breaker nodes, for a total of thirteen cluster and standby MDMs.

The following figure illustrates a 5-node MDM cluster:

Figure 1 5-node MDM cluster

All members of the MDM cluster have the same MDM package installed on them.

Before a server makes its way into the MDM cluster, it must follow the following path:

1. Install the MDM package on the server.
   During the installation, you determine if the server will be a Manager or a Tie Breaker (default).
2. Promote the server to Standby status, either as a Manager or as a Tie Breaker.
3. Add the standby server to the MDM cluster. A Manager, once entered into the cluster can take on the Master or Slave state.

MDM cluster creation is done automatically when deploying a system with any of the automated deployment tools.

Storage definitions

When configuring a ScaleIO system, you should take the following concepts into account: Protection Domains, Storage Pools, and Fault Sets. Together, these elements link the physical layer with the virtualization layer.

Protection Domains

A Protection Domain is a logical entity that contains a group of SDSs that provide backup for each other. Each SDS belongs to one (and only one) Protection Domain. Thus, by definition, each Protection Domain is a unique set of SDSs. In Figure 2 there are three Protection Domains. The one in the middle (fully depicted) consists of seven SDSs, each with two storage devices.

The maximum recommended number of nodes in a Protection Domain is 100. This enables the following:

- optimal performance
- reduction of theoretical *mean time between failure* issues
- ability to sustain multiple failures in different Protection Domains

You can add Protection Domains during installation. In addition, you can modify Protection Domains post-installation with all the management clients (except for OpenStack) .

Storage Pools

Storage Pools allow the generation of different storage tiers in the ScaleIO system. A Storage Pool is a set of physical storage devices in a Protection Domain. Each storage device belongs to one (and only one) Storage Pool. In Figure 2, there are 2 Storage Pools depicted.

When a volume is configured over the virtualization layer (see "SAN virtualization layer"), it is distributed over all devices residing in the same Storage Pool. Each volume block has two copies located on two different SDSs. This allows the system to maintain data availably following a single-point failure. The data will still be available following multiple failures, as long as each failure took place in a different storage pool.

Figure 2 Protection Domains and Storage Pools

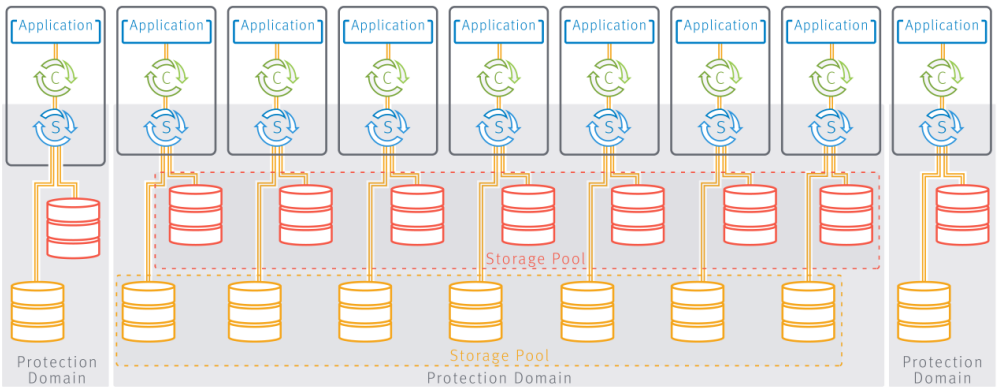

To provide consistent performance it is recommended that all devices in the Storage Pool will have similar storage properties.

For example, consider Figure 2. If all SDSs in a Protection Domain have two physical drives associated with them—one HDD and the other SSD— then you should define two Storage Pools:

- Capacity Storage Pool

  Consists of all HDDs in the Protection Domain
- Performance Pool

  Consists of all SSDs in the Protection Domain

Note

Mixing different types of media in the same pool is allowed, but be aware that due to the distribution of the data, performance will be limited to the least-performing member of the Storage Pool.

ScaleIO might not perform optimally if there are large differences between the sizes of the devices in the Storage Pool, for example, if one device is as big as the rest of the devices. If in doubt, contact ScaleIO support.

Each Storage Pool can work in one of the following modes:

- Zero padding enabled

  Ensures that every read from an area previously not written to returns zeros. Some applications might depend on this behavior. Furthermore, zero padding ensures that reading from a volume will not return information that was previously deleted from the volume.

  This behavior incurs some performance overhead on the first write to every area of the volume.
- Zero padding disabled (default)

  A read from an area previously not written to will return unknown content. This content might change on subsequent reads.

Zero padding must be enabled if you plan to use any other application that assumes that when reading from areas not written to before, the storage will return zeros or consistent data.

Note

The zero padding policy cannot be changed after the addition of the first device to a specific Storage Pool.

You can add Storage Pools during installation. In addition, you can modify Storage Pools post-installation with most of the management clients.

Fault Sets

A Fault Set is a logical entity that contains a group of SDSs within a Protection Domain, that have a higher chance of going down together, for example if they are all powered in the same rack. By grouping them into a Fault Set, you are telling ScaleIO that the data mirroring for all devices in this Fault Set, should take place on SDSs that are outside of this Fault Set.

When defining Fault Sets, we refer to the term fault units, where a fault unit can be either a Fault Set, or an SDS not associated with a Fault Set (you may think of it as a Fault Set of a single SDS).

There must be enough capacity within at least 3 fault units to enable mirroring.

If Fault Sets are defined, you can use any combination of fault units, for example:

- SDS1, SDS2, SDS3
- FS1, SDS1, SDS2
- FS1, FS2, SDS1
- FS1, FS2, FS3
  Figure 3 on page 39 illustrates the same configuration as Figure 2, with the addition of Fault Sets.

  Figure 3 Protection Domains, Storage Pools, and Fault Sets

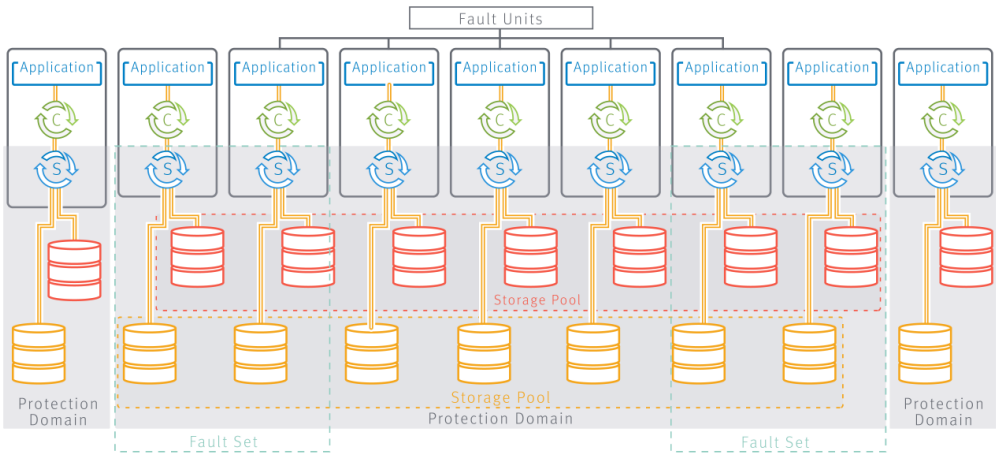

To use Fault Sets, you must work in the following order:

1. Ensure that a Protection Domain exists, or add a new one.
2. Ensure that a Storage Pool and Fault Sets (minimum of 3 fault units) exist, or add new ones.
3. Add the SDS, designating the PD and FS, and at the same time, adding the SDS devices into a Storage Pool.
   The automated deployment and installation tools follow this order automatically.

   You can only create and configure Fault Sets before adding SDSs to the system, and configuring them incorrectly may prevent the creation of volumes. An SDS can only be added to a Fault Set during the creation of the SDS.

   You define Fault Sets and add SDSs to them during installation, using the following management tools:

- Installation manager
- CLI
- REST
- VMware plug-in

In addition, you can also add Fault Sets when adding SDS nodes after initial installation.

Naming

It is recommended to name all ScaleIO objects with meaningful names. This will make it easier when defining volumes, associating them with applications, etc.

From the previous example, the Storage Pools can be named "Capacity_Storage" and "Performance_Storage," which allows you to identify the different tiers.

As for Protection Domains, one example would be separating the SDSs used by the finance department from those used by the engineering department. This segregation of different departments is very beneficial in many aspects (security being one of them). Thus, one might name the domains "Financial-PD" and "Engineering-PD."

The Fault Sets could be called "FS_Rack01" and "FS_Rack02."

Protection and load balancing

ScaleIO maintains the user data in a RAID-1 mesh mirrored layout. Each piece of data is stored on two different servers. The copies are randomly distributed over the storage devices. Rebuild and rebalance processes are fully automated, but are configurable.

Rebuild

When a failure occurs, such as on a server, device or network failure, ScaleIO immediately initiates a process of protecting the data. This process is called Rebuild, and comes in two flavors:

- Forward rebuild is the process of creating another copy of the data on a new server. In this process, all the devices in the Storage Pool work together, in a many-to-many fashion, to create new copies of all the failed storage blocks. This method ensures an extremely fast rebuild.
- Backward rebuild is the process of re-synchronization of one of the copies. This is done by passing to the copy only changes made to the data while this copy was inaccessible. This process minimizes the amount of data transferred over the network during recovery.

ScaleIO automatically selects the type of rebuild to perform. This implies that in some cases, more data will be transferred to minimize the time that the user data is not fully protected.

Rebuild throttling

Rebuild throttling sets the rebuild priority policy for a Storage Pool. The policy determines the priority between the rebuild I/O and the application IO when accessing SDS devices. Please note that application I/Os are continuously served.

Applying rebuild throttling will on one hand increase the time the system is exposed with a single copy of some of data, but on the other hand, will reduce the impact on the application. One has to make a decision and choose the right balance between the two.

The following possible priority policies may be applied:

- No Limit: No limit on rebuild I/Os.

  Any rebuild I/O is submitted to the device immediately, without further queuing. Please note that rebuild I/Os are relatively large and hence setting this policy will speed up the rebuild, but will have the maximal effect on the application I/O.
- Limit Concurrent I/O: Limit the number of concurrent rebuild I/Os per SDS device (default).

The rebuild I/Os are limited to a predefined number of concurrent I/Os. Once the limit is reached, the next incoming rebuild I/O waits until the completion of a currently executed rebuild I/O. This will complete the Rebuild quickly for best reliability, however, there is a risk of host application impact.

- Favor Application I/O: Limit rebuild in both bandwidth and concurrent I/Os.

  The rebuild I/Os are limited both in bandwidth and in the amount of concurrent I/Os. As long as the number of concurrent rebuild I/Os, and the bandwidth they consume, do not exceed the predefined limits, rebuild I/Os will be served. Once either threshold is reached, the rebuild I/Os wait until both I/O and bandwidth are below their thresholds. For example, setting the value to "1" will guarantee the device will only have one concurrent rebuild IO at any given moment, which will ensure the application IOs only wait for 1 rebuild IO at worst case.

  This imposes bandwidth on top of the Limit Concurrent I/Os option, which is a prerequisite to using this policy.
- Dynamic Bandwidth Throttling: This policy is similar to Favor Application I/O, but extends the interval in which application I/Os are considered to be flowing by defining a minimal quiet period. This quiet period is defined as a certain interval in which no application I/Os occurred. Note that the limits on the rebuild bandwidth and concurrent I/Os are still imposed.
- Default Values:
  - The default policy for rebuild is: Limit Concurrent I/O
  - Rebuild concurrent I/O Limit: 1 concurrent I/O

  Note

  Rebuild throttling affects the system's performance and should only be used by advanced users.

Rebalance

Rebalance is the process of moving one of the data copies to a different server. It occurs when ScaleIO detects that the user data is not evenly balanced across the devices in a Storage Pool. This can occur as a result of several conditions such as: SDS addition/removal, device addition/removal, or following a recovery from a failure. ScaleIO will move copies of the data from the most utilized devices to the least utilized ones.

Both Rebuild and Rebalance compete with the application IO for the system resources. This includes network, CPU and disks. ScaleIO provides a very rich set of parameters that can control this resource consumption. While the system is factory-tuned for balancing between speedy rebuild/rebalance and minimization of the effect on the application IO, the user has very fine-grain control over the rebuild and rebalance behavior.

Rebalance throttling

Rebalance throttling sets the rebalance priority policy for a Storage Pool. The policy determines the priority between the rebalance I/O and the application IO when accessing SDS devices. Please note that application I/Os are continuously served. Rebalance, unlike rebuild, does not impact the system's reliability and therefore reducing its impact is not risky.

The following possible priority policies may be applied:

- No Limit: No limit on rebalance I/Os.

  Any rebalance I/O is submitted to the device immediately, without further queuing. Please note that rebalance I/Os are relatively large and hence setting this policy will speed up the rebalance, but will have the maximal effect on the application I/O.

- Limit Concurrent I/O: Limit the number of concurrent rebalance I/Os per SDS device.

  The rebalance I/Os are limited to a predefined number of concurrent I/Os. Once the limit is reached, the next incoming rebalance I/O waits until the completion of a currently executed rebalance I/O. For example, setting the value to "1" will guarantee that the device will only have one rebalance IO at any given moment, which will ensure that the application IOs only wait for 1 rebalance IO in the worst case.

- Favor Application I/O: Limit rebalance in both bandwidth and concurrent I/Os.

  The rebalance I/Os are limited both in bandwidth and in the amount of concurrent I/Os. As long as the number of concurrent rebalance I/Os, and the bandwidth they consume, do not exceed the predefined limits, rebalance I/Os will be served. Once either limiter is reached, the rebalance I/Os wait until such time that the limits are not met again.

  This imposes a bandwidth limit on top of the Limit Concurrent I/Os option.

- Dynamic Bandwidth Throttling: This policy is similar to Favor Application I/O, but extends the interval in which application I/Os are considered to be flowing by defining a minimal quiet period. This quiet period is defined as a certain interval in which no application I/Os occurred. Note that the limits on the rebalance bandwidth and concurrent I/Os are still imposed.

- Default Values:
  - The default policy for rebalance: Favor Application I/O
  - Rebalance concurrent I/O Limit: 1 concurrent I/O per SDS device
  - Rebalance bandwidth limit: 10240 KB/s

Note

Rebalance throttling affects the system's performance and should only be used by advanced users.

Checksum protection

This feature addresses errors that change the payload during the transit through the ScaleIO system. ScaleIO protects data in-flight by calculating and validating the checksum value for the payload at both ends.

- During write operations, the checksum is calculated when the SDC receives the write request from the application. This checksum is validated just before each SDS writes the data on the storage device.
- During read operations, the checksum is calculated when the data is read from the SDS device, and is validated by the SDC before the data returns to the application. If the validating end detects a discrepancy, it will initiate a retry. The checksum will be done in the granularity of a sector (1/2KB).

This feature applies to all IOs: Application, Rebuild, Rebalance, and Migrate. The checksum is also kept in RMcache (Read Memory Cache), protecting every block that is maintained in SDS memory against memory corruption. The checksum feature can

be enabled at the Protection Domain level, and defined at the Storage Pool level. The feature is T10/DIF-ready.

Note

The checksum feature may have a major impact on performance and availability. Contact EMC customer support to verify if your use case is relevant.

ESX vStorage APIs for Array Integration (VAAI)

ESX vStorage APIs for Array Integration (VAAI) is a feature introduced in ESXi/ESX 4.1 that provides hardware acceleration functionality. It allows the host to offload specific virtual machine and storage management operations to compliant storage hardware. With the storage hardware's assistance, the host performs these operations faster, and consumes less CPU, memory, and storage fabric bandwidth.

VAAI uses these fundamental operations:

- Atomic Test & Set (ATS), which is used during creation and locking of files on the VMFS volume
- Clone Blocks/Full Copy/XCOPY, which is used to copy or migrate data within the same physical array
- Zero Blocks/Write Same, which is used to zero-out disk regions
- Thin Provisioning in ESXi 5.x and later hosts, which allows the ESXi host to tell the array when the space previously occupied by a virtual machine (whether it is deleted or migrated to another datastore) can be reclaimed on thin provisioned LUNs.
- Block Delete in ESXi 5.x and later hosts, which allows for space to be reclaimed using the SCSI UNMAP feature.

The ScaleIO supported VAAI features are:

- Atomic Test & Set (ATS)
- Zero Blocks/Write Same
- Thin Provisioning in ESXi 5.x and later hosts
- Block Delete in ESXi 5.x and later hosts

The following output is an example of typical output:

```
esxcli storage core device vaai status get -d
eui.7dbf14034834bbe01bf7e55800000002
eui.7dbf14034834bbe01bf7e55800000002
VAAI Plugin Name:
ATS Status: supported
```

`Clone Status: unsupported` This means that Clone Block/Full Copy/Xcopy is not supported.

`Zero Status: supported` This means that write same is supported.

`Delete Status: supported` This means that UNMAP is supported.

Note

Thin provisioning is not shown in VAAI output.

Caching

ScaleIO offers a number of caching options, for the purpose of enhancing system performance.

The following caching options are supported by ScaleIO:

- RAM Read Cache (using DRAM server memory)
- Read Flash Cache (using SSD devices)

In addition, the following caching solutions are available:

- CacheCade (using SSD devices) - available in VxRack Node 100 Series systems
- DAS Cache (using SSD devices) - available in ScaleIO Ready Node systems

Note

DAS Cache is not supported by PowerEdge R640, R640xd, or R740xd servers.

SSDs used for caching cannot be used for storage purposes.

The following table summarizes information about the caching modes provided by the system.

Table 5 Caching modes

Mode	Description	Considerations	Default Setting
RAM Read Cache (RMcache)	Read-only caching performed by server RAM.	RAM Read cache, the fastest type of caching, uses RAM that is allocated for caching. Its size is limited to the amount of allocated RAM. **Note** The amount that may be allocated is limited, and can never be the maximum available RAM.	Disabled, except when storage-only nodes are deployed.
RFcache	Read-only caching performed by one or more dedicated SSD devices in the server.	RFcache uses the full capacity of SSD devices (up to eight) to provide a larger footprint of read-only LRU (Least Recently Used) based-caching resources for the SDS. This type of caching reacts quickly to workload changes to speed up HDD Read performance. Several SSD devices can be allocated to a shared cache pool, and therefore the cache size is limited in size only by the amount of SSDs allocated for this purpose. The RFcache driver must be installed during deployment. Caching devices can be defined either during the installation process or after deployment. Limitations: RFcache does not support partitions on devices installed on Windows nodes.	

Table 5 Caching modes (continued)

Mode	Description	Considerations	Default Setting
		Support matrix: • An RFcache device (flash device) can be partitioned only on Linux. • An SDS storage/source device cannot be partitioned if it needs to be accelerated by RFcache. • An SDS storage/source device as a file (over file system), cannot be accelerated by RFcache.	
CacheCade	Read and write-back caching performed by one or more dedicated SSD devices in the server.	CacheCade uses the full capacity of one or more SSD devices to provide a large footprint of both read and write-back caching resources to the SDS. This caching mode moves "hot" (active) chunks of data from HDDs to cache, for Read and Write buffering. For write-back caching, the write is temporarily written to the SSD, which is much faster than an HDD, allowing faster response of the SDS to write acknowledgment. Two SSD devices can be allocated to a shared cache pool, up to a maximum size of 512 GB in total. **Note** If a fault occurs in the caching device before the writes have been offloaded, all the HDD devices cached by CacheCade acquire failed status, and a rebuild process commences in VxRack Node 100 Series. Once the rebuild is over, the caching disk can be replaced, all caching has stopped in the storage pool, and the HDD members in the storage pool can be cleared of errors.	Disabled
DAS Cache	Read and write-back caching performed by one or more dedicated SSD devices in the server	DAS Cache uses the full capacity of one or more SSD devices to provide a large footprint of both read and write-back caching resources to the SDS. This caching mode moves "hot" (active) chunks of data from HDDs to cache, for Read and Write buffering. For write-back caching, the write is temporarily written to the SSD, which is much faster than an HDD, allowing faster response of the SDS to write acknowledgment. One SSD device can accelerate several HDDs (in DAS Cache they are called "Volumes"). Striping the Cache on two devices is not supported in the ScaleIO Ready Node solution.	Disabled

Table 5 Caching modes (continued)

Mode	Description	Considerations	Default Setting
		Note If a fault occurs in the caching device before the writes have been offloaded, all the HDD devices cached by DAS Cache acquire failed status, and a rebuild process commences in ScaleIO. Once the rebuild is over, the caching disk can be replaced, all caching has stopped in the storage pool, and the HDD members in the storage pool can be cleared of errors.	

The following table illustrates the caching support matrix:

Table 6 Caching support matrix

System	RFcache	RMcache	DAS Cache	CacheCade
ScaleIO	Yes	Yes		
ScaleIO Ready Node	Yes	Yes	Yes	
VxRack Node 100 Series	Yes	Yes		Yes

Networking

In ScaleIO, inter-node communication (for the purposes of managing data locations, rebuild and rebalance, and for application access to stored data) can be done on one IP network, or on separate IP networks. Management (via any of the management interfaces) can be done in the following ways:

- Via a separate network with access to the other ScaleIO components
- On the same network

These options can be configured a) during deployment in the full Installation Manager (via the CSV topology file) and using the VMware plug-in, as well as b) after deployment with the CLI.

This section describes how to choose from these options, depending on your organization's requirements, security considerations, performance needs, and IT environment.

ScaleIO networking considerations:

- Single IP network: All communications and IOs used for management and for data storage are performed on the same IP network. This setup offers the following benefits:
  - Ease of use
  - Fewer IP addresses required
- Multiple separate IP networks: Separate networks are used for management and for data storage, or separate networks are used within the data storage part of the system. This setup offers the following benefits:
  - Security

- Redundancy
- Performance
- Separate IP roles in order to separate between customer data and internal management

Note

Network high availability can be implemented by using NIC-bonding (refer to relevant operating system vendor guidelines for best practices) or by using several data networks in ScaleIO.

For more information about MTU performance considerations and best practices, see the *ScaleIO Performance Fine-Tuning Technical Notes*.

The following table describes the range of potential IP address configurations:

Table 7 IP address configurations in ScaleIO (based on CSV file)

Column in CSV file	MDM Mgmt IP	MDM IPs	SDS All IPs	SDS-SDS Only IPs	SDS-SDC Only IPs
Comments	Management Access	Control Network	Rebuild and Data Path Network	Rebuild Network	Data Path Network
	Optional, but recommended; not applicable for Tie Breaker IP addresses that can be used to provide access to ScaleIO management applications, such as CLI, GUI, REST API, OpenStack. This IP address must be externally accessible.	Mandatory IP addresses used for MDM control communications with SDSs and SDCs, used to convey data migration decisions, but no user data passes through the MDM. Must be on the same network as the data network. Must be externally accessible if no MDM Management IP addresses are used.	IP addresses used for both SDS-SDS and SDS-SDC communications. These IP addresses will also be used to communicate with the MDM	IP addresses used for SDS-SDS communication only. These addresses are used for rebuild & rebalance operations. These IP addresses will also be used to communicate with the MDM.	IP addresses used for SDS-SDC communication. These addresses are only used for read-write user data operations.

The following combinations can be used for SDS/SDC:

- Only *SDS All IPs*
- Only *SDS-SDS Only IPs* + *SDS-SDC Only IPs*
- *SDS All IPs* + either *SDS-SDS Only IPs* or *SDS-SDC Only IPs* (can be used in cases of multiple networks; ensure that you do not use the same IP address more than once in the networks).
- *SDS All IPs* + both *SDS-SDS Only IPs* and *SDS-SDC Only IPs* (can be used in cases of multiple networks; ensure that you do not use the same IP address more than once in the networks).

Note

On Linux nodes, only the MDM needs a management IP address.

On Windows nodes, only the MDM needs a management IP address.

On VMware, all ScaleIO VMs need to have a management IP address as well as another address for the data network, the network on which traffic flows between SDSs and SDCs for read/writes, rebuild, and rebalance.

In the following example drawing for separate networks, a very simple example is shown, where the management and storage parts of the system are on different networks. In more complex configurations, MDMs, SDCs and SDSs can be on separate networks. Up to 8 separate networks per ScaleIO system are supported.

The following figures show example configurations and the corresponding fields in a CSV configuration file:

Figure 4 ScaleIO system deployed on a single network

Example of a ScaleIO system deployed on a single net

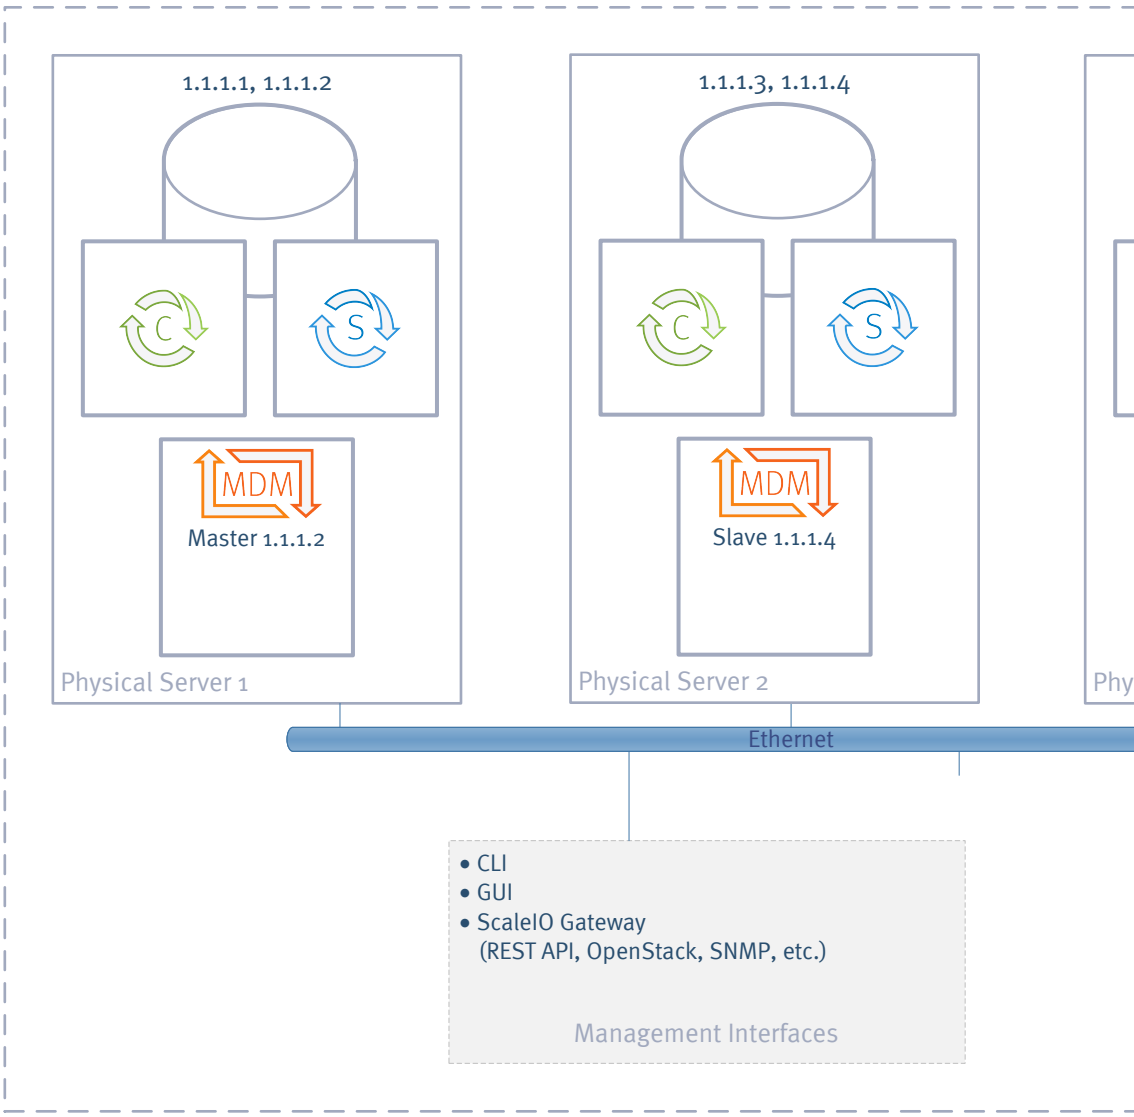

	A	B	C	D	E	F	G	H
1	Username	Password	Operating System	Is MDM/TB	MDM Mgmt IP	MDM IPs	Is SDS	SDS Name
2	root	Password1	linux	Master		1.1.1.2	Yes	SDS1
3	root	Password1	linux	Slave		1.1.1.4	Yes	SDS2
4	root	Password1	linux	TB			Yes	SDS3

Figure 5 ScaleIO system deployed on separate networks

Example of ScaleIO storage and management deployed on sep

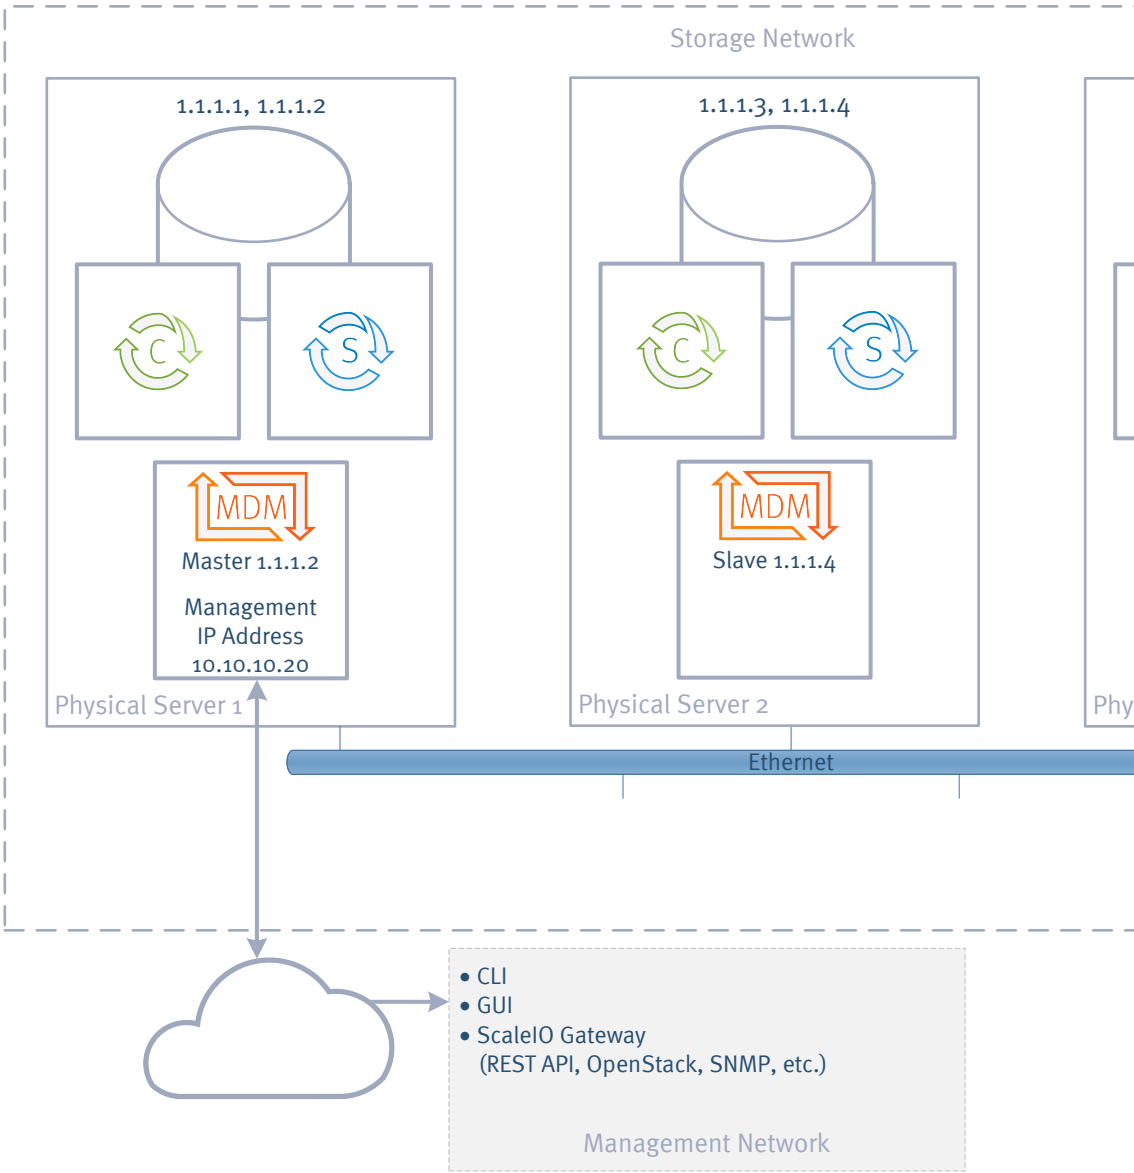

	A	B	C	D	E	F	G	H	
1	Username	Password	Operating System	Is MDM/TB	MDM Mgmt IP	MDM IPs	Is SDS	SDS Name	SI
2	root	Password1	linux	Master	10.10.10.20.	1.1.1.2	Yes	SDS1	1.
3	root	Password1	linux	Slave		1.1.1.4	Yes	SDS2	1.
4	root	Password1	linux	TB			Yes	SDS3	1.

VMware limitation:

Multiple IP subnets used for the ScaleIO Data network cannot be on the same subnet in a VMware setup.

For more information, see the VMware limitation in the following link:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2010877

ScaleIO only supports the following network configurations when deployed on VMware:

- A single data storage network
- Two or more data networks, each on separate IP subnets
- A single IP data network using several NIC-bonding configurations, or vSwitch load balancing

Virtual IP Address

Up to four virtual IP addresses can be defined for the MDM cluster. SDCs are then mapped to the MDM cluster's virtual IP addresses, instead of to static MDM IP addresses. MDMs are sometimes switched during normal operation of the cluster, and the virtual IP address will always be mapped to the active MDM. The use of virtual IP addresses simplifies maintenance procedures on the MDM cluster, because system components communicate via the virtual IP addresses. Therefore, SDCs do not need to be reconfigured when a server hosting an MDM is replaced.

Note

Virtual IP addresses are not currently supported on Windows-based systems.

In new installations in Linux environments, the MDM cluster's virtual IP address can be added and mapped using the Installation Manager CSV file. In VMware environments, virtual IP addresses are mandatory, and configuration is performed using the ScaleIO VMware Installation Wizard, in the Configure SVM stage. The REST API can also be used to add virtual IP addresses to the cluster. In all cases, a virtual IP NIC placeholder must be mapped to each virtual IP address. Ensure that there are NICs available for this purpose.

Existing systems may be extended to include additional MDMs to a cluster. The new MDMs should be mapped to the existing virtual IP addresses.

If virtual IP addresses need to be modified, you must use the CLI or the REST API (not the IM or the vSphere plug-in), and it must be done with extreme caution.

All SDCs will require reconfiguration, to reflect the changes made to the MDM cluster. Otherwise, the SDCs will not be able to communicate with the MDM cluster, and volumes will not be accessible.

S.M.A.R.T. hardware monitoring

The ScaleIO bare-metal solution now provides monitoring capabilities for RAID controllers and storage devices compatible with S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) protocols.

In Linux-based environments, S.M.A.R.T.-compatible HDDs, SSDs and RAID storage controllers can be monitored for S.M.A.R.T. attributes such as temperature, SSD wear level, and error counters. LEDs can also be lit on these hardware devices, to simplify physical identification for maintenance purposes.

Each hardware vendor defines specific thresholds for the S.M.A.R.T. attributes. This feature currently supports storage devices controlled by LSI, HP and Dell RAID controllers, and stand-alone devices. During system deployment, an external

monitoring tool is installed as part of the LIA on each node. Additional RAID controller tools must be installed manually after system deployment: storcli for LSI RAID controllers, hpssacli for HP RAID controllers, or perccli for DELL RAID controllers. These tools are used by the system to collect the counters that are returned to the MDM.

Note

In some cases, LSI RAID controllers may report vendor information as "AVAGO" instead of LSI.

The MDM queries the SDSs at set intervals, and stores the returned information. This information can be viewed using CLI queries. In addition, when thresholds are crossed for S.M.A.R.T. attributes, alerts are generated by the system.

When the CLI is used to query device information, physical device information, such as serial number, model name, vendor etc., temperature, and wear level information (for SSDs only) is included in the returned response.

For information about the use of CLI commands, see the *ScaleIO CLI Reference Guide*.

For information about the use of REST API URIs, see the *ScaleIO User Guide*.

You can use the GUI to monitor S.M.A.R.T.-related alerts in the **Alerts** view.

In addition, SNMP traps and ESRS alert codes can be used to monitor alerts triggered by devices compatible with S.M.A.R.T.

List of approved RAID controllers

Provides high-level specifications of RAID controllers, which are tested and certified by ScaleIO.

ScaleIO-certified RAID controllers

The following table describes the ScaleIO-certified RAID controllers:

Manufacturer	Specifications
HP	• Model Name: Smart Array P440ar • Vendor Name: HP • Firmware Version: 3.56 • Driver Version: 3.4.10 • Driver Name: hpsa • PCI Address: 0000:03:00.0
DELL	• Model Name: PERC H730 Mini • Vendor Name: Dell • Firmware Version: 25.3.0.0016 • Driver Version: 06.807.10.00-rh • Driver Name: megaraid_sas • PCI Address: 00:02:00:00
LSI	• Model Name: LSI MegaRAID SAS 9271-8i • Vendor Name: LSI • Firmware Version: 23.12.0-0021

Manufacturer	Specifications
	• Driver Version: 06.810.09.00-rh • Driver Name: megaraid_sas • PCI Address: 00:82:00:00
	• Model Name: LSI MegaRAID SAS 9271-8i • Vendor Name: LSI • Firmware Version: 23.12.0-0018 • Driver Version: 06.805.06.01-rc • Driver Name: megaraid_sas • PCI Address: 00:82:00:00
	• Model Name: LSI MegaRAID SAS 9271-8i • Vendor Name: LSI • Firmware Version: 23.12.0-0021 • Driver Version: 06.805.06.01-rc • Driver Name: megaraid_sas • PCI Address: 00:82:00:00
	• Model Name: LSI MegaRAID SAS 9271-8i • Vendor Name: LSI • Firmware Version: 23.34.0-0005 • Driver Version: 06.810.08.00 • Driver Name: megaraid_sas • PCI Address: 00:82:00:00

Snapshots

The ScaleIO storage system enables you to take snapshots of existing volumes, up to 31 per volume. The snapshots are thin provisioned and are extremely quick. For more information about thin provisioning, see "SAN virtualization layer".

Once a snapshot is generated, it becomes a new, unmapped volume in the system. You can manipulate it in the same manner as any other volume exposed by the ScaleIO storage system.

Figure 6 Snapshot operations

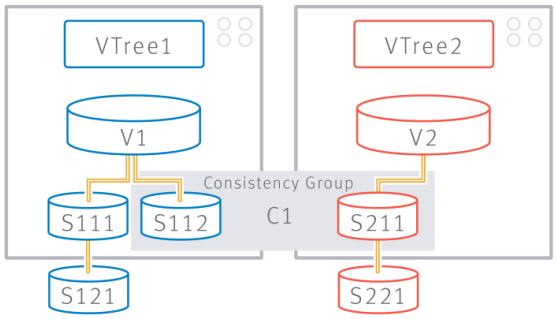

The structure related to all the snapshots resulting from one volume is referred to as a VTree (short for Volume Tree). It is a tree spanning from the source volume as the root, whose siblings are either snapshots of the volume itself or descendants of it. Thus, some snapshot operations are related to the VTree, and may affect parts of it. In Figure 8 in BLUE, S111 and S_{112} are snapshots of V1. S121 is a snapshot of snapshot S111. Together, V1 and S1xy are the VTree of V1.

When taking a snapshot in the system, you can specify more than one volume. All snapshots taken together form a consistency group. They are consistent in the sense that they were all taken at the same time. So if there is a contextual relationship between the data contained by all the snapshot members, then that set is meaningful. The consistency group allows manipulation of the entire set.

If you remove an entire consistency group, all of the snapshots that were taken together will be removed. Back to Figure 8, in RED, S211 is a snapshot of V2. Since S112 and S211 were taken together, they compose a consistency group (in ORANGE) designated as C1.

Note

The consistency group is only for convenience purposes. There are no protection measures done by ScaleIO to conserve the consistency group. For example, you can remove a snapshot that is a member of a consistency group.

V-Trees

A V-Tree (short for volume tree) is the structure comprised of a volume and the snapshots resulting from that volume. It is a tree spanning from the source volume at its root, whose descendants are either snapshots of the volume itself or snapshots of a snapshot. In the V-Tree diagram, S_{111} and S_{112} are snapshots of V_1. S_{121} is a snapshot of snapshot S_{111}. Together, V_1 and S_{1xy} are the V-Tree of V_1.

Figure 7 V-Tree diagram

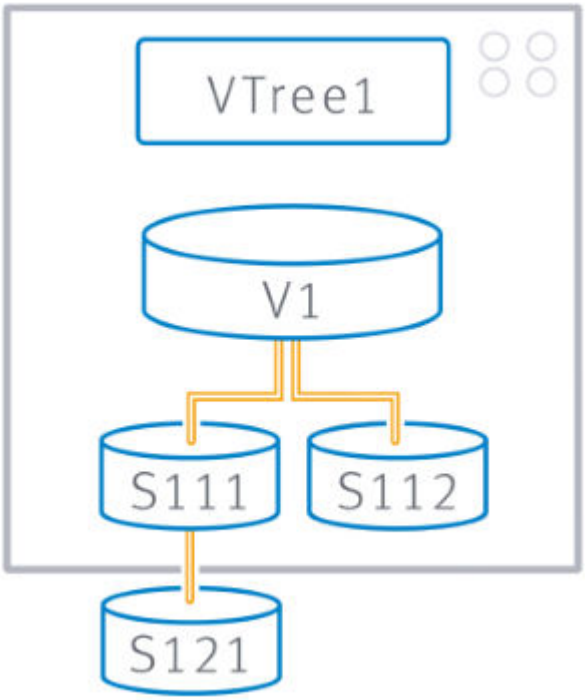

Other functions

ScaleIO includes the following functions:

- EMC Secure Remote Support (ESRS)

  ESRS support enables secure, high-speed, 24x7, remote connection between EMC and customer installations, including:

  - Remote monitoring
  - Remote diagnosis and repair
  - Daily sending of logs, alerts, and ScaleIO topology

- Syslog

  The MDM syslog service can send events, via TCP/IP, to RFC 6587-compliant remote (or local) Syslog servers. Messages are sent with facility local0, by default. Once the syslog service is started, all events will be sent until the service is stopped.

- Get Info

  Get Info assembles a ZIP file of system logs for troubleshooting. You can run this function from a local node for its own logs, using the CLI, or by using the Installation Manager to assemble logs from all MDM and SDS nodes in the system. In addition to the log files, a visual snapshot of the ScaleIO GUI, from the time you perform the operation, can be saved, to better enable support options.

  The Get Info function is described in the *Log Collection Guide*.

- Quality of Service (QoS)

You can adjust the amount of bandwidth and storage that any given SDC can use. You can configure this with the CLI and the REST interface, on a per client/per volume basis.

- Background Device Scanner

  The Background Device Scanner ("scanner") enhances the resilience of your ScaleIO system by constantly searching for, and fixing, device errors before they can affect your system. This provides increased data reliability than the media's checksum scheme provides. The scanner seeks out corrupted sectors of the devices in that pool, provides SNMP reporting about errors found, and keeps statistics about its operation.

  When a scan is completed, the process starts again, thus adding constant protection to your system.

  You can set the scan rate (default: 1 MB/second per device), which limits the bandwidth allowed for scanning, and choose from the following scan modes:

  - Device only mode

    The scanner uses the device's internal checksum mechanism to validate the primary and secondary data. If a read succeeds in both devices, no action is taken. If a faulty area is read, an error will be generated.

    If a read fails on one device, the scanner attempts to correct the faulty device with the data from the good device. If the fix succeeds, the error-fixes counter is increased. If the fix fails, a device error is issued.

    Note

    A similar algorithm is performed every time an application read fails on the primary device.

    If the read fails on both devices, the scanner skips to the next storage block.

  - Data comparison mode (only available if zero padding is enabled)

    The scanner performs the same algorithm as above, with the following additions:

    After successful reads of the primary and secondary copies of the data, the scanner calculates and compares their checksums. If this comparison fails, the compare errors counter is increased, and the scanner attempts to overwrite the secondary device with the data from the primary device. If this fails, a device error is issued.

    The scanning function is enabled and disabled (default) at the Storage Pool level, and this setting affects all devices in the Storage Pool. You can make these changes at anytime, and you can add/remove volumes and devices while the scanner is enabled.

    When adding a device to a Storage Pool in which the scanner is enabled, the scanning will start about 30 seconds after the device is added.

- AD over LDAP or LDAPS authentication

  User authentication may be done using AD (Active Directory) over LDAP (Lightweight Directory Access Protocol) or LDAPS (Secure LDAP). ScaleIO can support both AD users that are fully controlled through the customer’s existing centralized location, and local users (as has been supported in earlier ScaleIO versions). You can associate groups from the AD with the existing ScaleIO roles in order to ensure the Role-Based Access (RBAC) model. When a user logs on to the

ScaleIO system, the MDM identifies that the user belongs to the AD domain, and authenticates the user against the AD server over secured communications. Once the user is authenticated, ScaleIO accepts the group to which the user belongs according to the AD, and associates the appropriate role and its user permissions to that user. The AD implementation is fully redundant.

Note

The authorization permissions of each role are defined differently for local authentication, and for LDAP/LDAPS authentication.

The benefits of using AD over LDAP/LDAPS include:

- Full control of ScaleIO users through the main user repository
- No need to specify a local user for each customer

  If the AD directory is down, the administrator can always use local users to maintain the ScaleIO system.

- Oscillating failure handling

  The Oscillating Failures feature detects and reports various oscillating failures, in cases when components fail repeatedly and cause unnecessary failovers, and therefore disruptions to normal system operation. Typical examples of oscillating failures include:

  - A disk that accepts some I/Os and rejects others
  - A node with interrupted connectivity
  - A node that is constantly busy and therefore handles some I/Os too slowly
  - A disk that is sometimes slow to respond
  - A network that is experiencing disruptions

    The smart detection of such failures provides the ability to handle error situations, and to reduce their impact on normal system operation. Oscillating failure handling can be set for MDMs, SDSs and for SDCs. For SDSs, failure handling can be defined per Protection Domain or per Storage Pool.

  - Oscillating failure counters

    The following table describes the oscillating failure counters:

Oscillating failure counters	Description
(sds_sds/sdc_mdm/sdc_sds/mdm_sds) network_disconnections	Measures the number of network disconnections (socket closed) between two components per IP address
sds_decoupled	Measures the number of times an SDS process is down, as detected by the MDM
sds_configuration_failures	Measures the number of times the MDM fails to configure an SDS, when connecting to an SDS (failures occur during the reconfiguration phase)

Oscillating failure counters	Description
sds_receive_buffer_allocation_failures	Measures the number of times an SDS fails to allocate buffer for receiving messages
sdc_long_operations	Measures the number of SDC RPC operations that take longer than the predefined threshold (default threshold is 5 seconds)
sdc_memory_allocation_failures	Measures the number of memory allocation failures in each SDC
sdc_socket_allocation_failures	Measures the number of socket allocation failures in each SDC
sds_device_long_successful_ios	Measures the number of successful IOs to an SDS device, which take longer than the predefined threshold (default threshold is 250 milliseconds)

- Secure connectivity with external components

  This feature allows external components to authenticate the MDM. After authentication, communication between the MDM and external components is performed using TLS (Transport Layer Security) protocols.

  Secure communication with the MDM is authenticated by the following components:

  - CLI client

    Note

    If the secure mode is not enabled, modifications are necessary to run SCLI commands.

  - ScaleIO Gateway
  - GUI client
  - IM client

    Once added in the trust point, all communications will require authentication, followed by communications over TLS. The same method is employed between the IM and all LIAs.

Implementing ScaleIO

Implementing a ScaleIO system is, in general, a two-step process: first build the physical storage layer, then configure the virtual SAN layer on top of it.

Physical layer

The physical layer consists of the hardware (servers with storage devices and the network between them) and the ScaleIO software installed on them.

Typically, each SDS is physically located on a separate server, but ScaleIO also supports the installation of multiple SDSs per server.

The multiple SDS feature allows you to take fuller advantage of the server's computing resources, particularly when no applications are running along-side the SDS. Each SDS on the server is unique, with its own name (for example, sds1, sds2), path, and ports used. This uniqueness provides better control over the SDSs installed. Each SDS can be installed, removed, or upgraded independently. It is up to the user to control the scope of the SDS object. For example, if the SDSs are placed in the same Protection Domain, the user must put them in a single Fault Set.

ScaleIO currently supports four SDSs running on each server. Each SDS is installed using a different RPM. The user can start with a single SDS, and add more SDSs later.

Currently, there is no Windows support for the Multiple SDS feature.

To implement the physical layer, perform the following steps:

1. Install the MDM component on the MDM nodes in one of the following configurations:
   - Three-node redundant cluster (one Master MDM, one Slave MDM, and one Tie Breaker).
   - Five-node redundant cluster (one Master MDM, two Slave MDMs, and two Tie Breakers).
   - Single node (one master MDM).

     NOTICE

     It is not recommended to use Single Mode in production systems, except in temporary situations. The MDM contains all the metadata required for system operation. Single Mode has no protection, and exposes the system to a single point of failure.

     MDMs do not require dedicated nodes. They can be installed on nodes hosting other ScaleIO components.
2. Install the SDS component on all nodes that will contribute some, or all, of their physical storage.

   Divide the SDS nodes into Protection Domains. Each SDS can be a member of only one Protection Domain.

   Per Protection Domain, divide the physical storage units into Storage Pools, and optionally, into Fault Sets.
3. Install the SDC component on all nodes on which the application will access the data exposed by the ScaleIO volumes.

Figure 8 Physical layout example—3-node cluster

Communication is done over the existing LAN using standard TCP/IP. The MDM and SDS nodes can be assigned up to eight IP addresses, enabling wider bandwidth and better I/O performance and redundancy.

You can perform physical layer setup using the following methods:

- ScaleIO Installation Manager a web-client based tool
- ScaleIO VMware plug-in a VMware plug-in
- Manual installation procedures

  After completing this installation, the physical layer is ready, and can expose a virtual storage layer.

SAN virtualization layer

The MDM cluster manages the entire system. It aggregates the entire storage exposed to it by all the SDSs to generate a virtual layer - virtual SAN storage. Volumes can now be defined over the Storage Pools and can be exposed to the applications as a local storage device using the SDCs.

To expose the virtual SAN devices to your servers (the ones on which you installed and configured SDCs), perform the following:

- Define volumes. Each volume defined over a Storage Pool is evenly distributed over all members using a RAID protection scheme. By having all SDS members of the Storage Pool participate, ScaleIO ensures:
  - Highest and most stable and consistent performance possible
  - Rapid recovery and redistribution of data
  - Massive IOPS and throughput

    You can define volumes as follows:

- Thick

  Capacity is allocated immediately, even if not actually used. This can cause capacity to be allocated, but never used, leading to wasted capacity.

  Thick capacity provisioning is limited to available capacity.

- Thin

  Capacity is "on reserve," but not allocated until actually used. This policy enables more flexibility in provisioning.

  Whereas thick capacity is limited to available capacity, thin capacity provisioning can be oversubscribed, as follows:

  Maximum thin capacity provisioning = 5 * (gross capacity - used capacity)

  When capacity usage reaches the level where it may cause IO errors, alerts are generated. At certain higher capacity levels, volumes (even thin volumes) can no longer be created.

  Example:

  In a system with 3 SDSs, each with 10 TB, there are 30 TB of storage.

  In the system, there is already a thick-provisioned volume that takes up 15 TB of the gross capacity (created by adding a 7.5 TB volume).

  MDM will allow a total of 300 TB gross to be provisioned, and since 15 TB are already allocated, you can add a thin-provisioned volume of 285 TB gross (by adding a 142.5 TB volume) or a thick-provisioned volume of 15 TB gross.

- Map volumes. Designate which SDCs can access the given volumes. This gives rise to the following:
  - Access control per volume exposed
  - Shared nothing or shared everything volumes

    Once an SDC is mapped to a volume, it immediately gets access to the volume and exposes it locally to the applications as a standard block device. These block devices appear as `/dev/sciniX` where `X` is a letter, starting from "a."

    For example:

    - `/dev/scinia`
    - `/dev/scinib`
  - The maximum amount of partitions for the scini disk is 15.
- In a Windows environment, the device looks like any other local disk device, as shown in the Device Manager.

  The maximum amount of volumes that can be mapped to an SDC is listed in the "Product limits" table.

Note

SDC mapping is similar to LUN mapping, in the sense that it only allows volume access to clients that were explicitly mapped to the volume.

This is the end of the system setup.

Implementing ScaleIO over a virtual system

This section provides an overview of how ScaleIO is implemented in a virtualized environment.

VMware implementation

ScaleIO deployment in the VMware environment is implemented differently.

Implementation
In the VMware environment, the MDM and SDS components are installed on a dedicated SVM, whereas the SDC is installed directly on the ESX host.

Note

Installing the SDC on the ESX host requires a restart of the ESX host.

This implementation is illustrated in the following figure:

Figure 9 ScaleIO implementation on ESX

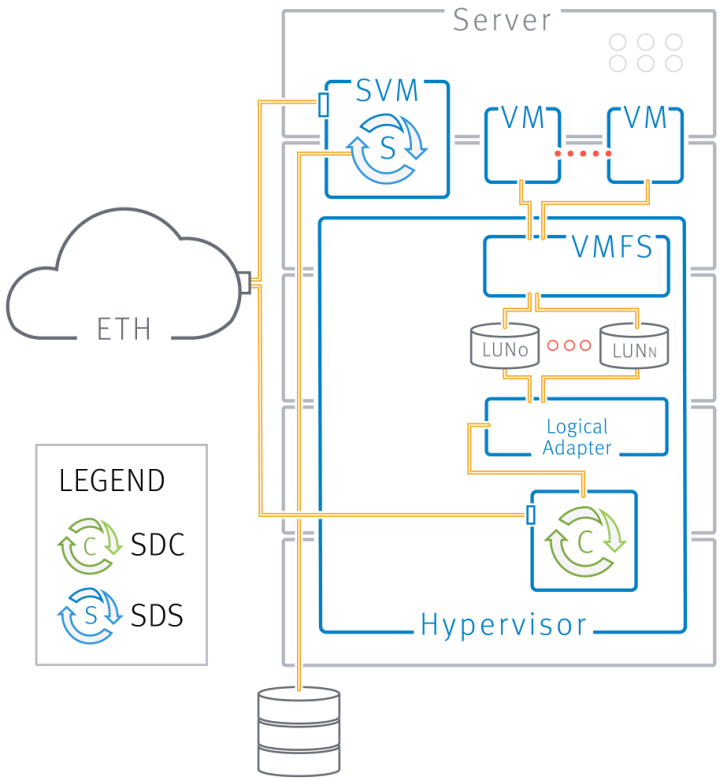

The LUNs in the previous figure can be formatted with VMFS, and then exposed using the ESXi host to the virtual machine, or can be used as RDM devices. When the LUNs are used as RDM devices, the VMFS layer is omitted.

Installation in a VMware environment is enabled via the VMware plug-in.

Device management
In the VMware environment, devices can be managed in the following ways:

- VMDirectPath I/O
  Device management is performed via the SVM, yielding the best I/O performance. Devices are added to the system after the deployment. On factory-installed ScaleIO Ready Node servers, the vSphere plug-in configures DirectPath on the ESX servers and adds devices to the system. On other servers, you must enable VMDirectPath manually to the ESX server, then use the plug-in to add devices.

  Requirements:

  - When DirectPath is configured, all servers in the system must use DirectPath. For ScaleIO Ready Node or VxRack Node 100 Series systems, this can be configured with the ScaleIO vSphere plug-in. For other servers, use the vSphere client to configure each ESX host manually.
  - If the host has multiple controllers, you must configure DirectPath on that host manually, not with the plug-in. After that is done, you can use the plug-in to deploy ScaleIO on that host and to add devices to it.
  - All devices on all servers must not have any other use (not VMDK, RDM, or be part of a datastore).
  - Before using the vSphere plug-in to add devices, ensure that the devices have VD names. To add unnamed devices, use the SCLI `add_sds` command. For more information, see the "Add SDS device with vSphere plug-in fails in DirectPath environment" troubleshooting section in the *ScaleIO User Guide*.
  - PowerEdge 14G servers with any NVME devices cannot be added in a DirectPath-based system. Use RDM-based, instead.
  - ESX boot device requirements:
    - Must be on a separate controller or connected directly to the board.
    - SATADOM and M2 boot devices are supported. These devices allow the creation of a datastore on the system disk, which is needed to host the ScaleIO VM.
    - USB boot devices are not supported.
- RDM

  Using RDM mapping, a device is created on the SVM that points to the physical disk on the ESX.

  You can add RDM devices that are connected through a physical RAID controller. If a local RDM is not connected via a RAID controller, it may not be supported. To ensure the compatibility of these devices, you can add them as VMDK, or you can select **Enable RDMs on non parallel SCSI controllers**, as described in the "Advanced settings options" section of the *ScaleIO Deployment Guide*. Enable this option before beginning the deployment.

  Before enabling this feature, contact EMC Support.
- VMDK

  A new datastore is created, with a VMDK, and the VMDK is added to the SVM. ScaleIO requires thick provisioning, so this process can take a long time.

  In almost all cases, RDM is the preferred method to add physical devices. Use the VMDK method only in the following scenarios:

  - The physical device does not support RDM.
  - The device already has a datastore, and the device isn't being completely used. The excess area that is not already being used will be added as a ScaleIO device.

Note

To use VMDK, select **Enable VMDK creation**, as described in the "Advanced settings options" section of the *ScaleIO Deployment Guide*.

System size

If you are deploying a very large ScaleIO system (several hundred nodes), you can increase the parallelism limit (default: 100), thus speeding up the deployment. This is dependent on the processing power of the vCenter.

To increase the parallelism limit, use the plug-in **Advanced settings**, as described in the "Advanced settings options" section of the *ScaleIO Deployment Guide*.

Pre-deployment considerations

You should take these considerations into account before deploying the system:

- Do you want to use separate networks for data and management (recommended), and which IP addresses will you use for the SVMs and VMkernels?
- Are there flash devices that will be added to the SDS?
- Do you want to create Fault Sets? See the requirements in Fault Sets on page 38.

Post-deployment considerations

You should take these considerations into account after deploying the system:

- After deployment is complete, set all SVMs to start automatically with the system. Do not set SVMs under the VMware resource-pool feature.
- In a DirectPath environment, after deploying the system, you must add devices to the SDS.

Xen implementation

In a Xen environment, both the SDC and SDS are installed in Dom0 as would be on a physical node. Dom0 accesses the storage media through the SDS and exposes volumes based on ScaleIO through the SDC.

Figure 10 ScaleIO Xen virtual machine architecture

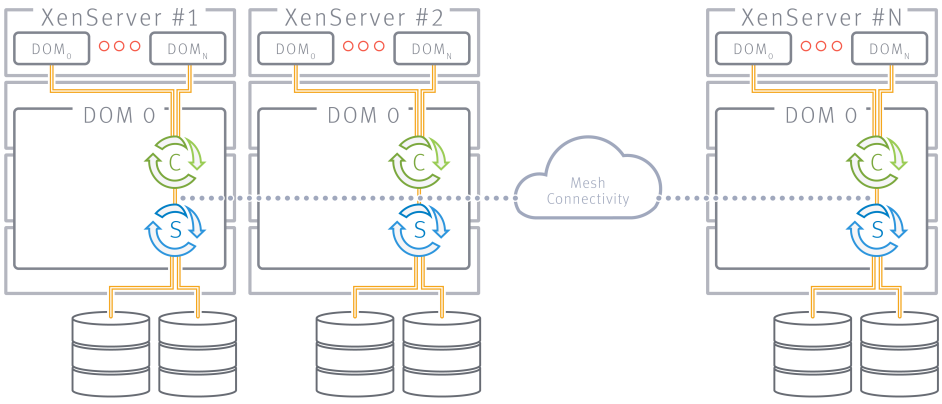

Information on provisioning in a Xen environment is described in this guide.

Maintenance

Maintenance of ScaleIO is primarily limited to configuration changes of the physical and virtual layers. It requires minimal user attention. When maintenance or planned

restart of an SDS is required, the maintenance mode feature can be used to streamline system operation.

Maintaining the physical layer

In the physical layer, maintenance is limited to adding and removing hardware units and configuring them into the ScaleIO system. These operations are usually a result of:

- Scaling out when there is a need for additional capacity. This usually results in adding more storage media to the existing servers, or adding additional servers.
- Hardware failure. In cases where there is a hardware (storage media or server) failure and it needs to be replaced.

  In all of the above cases, the operation will require adding or removing storage capacity from the system. In some cases, it may include adding or removing an entire server, and its associated storage media, from the configuration. As far as ScaleIO is concerned, all of these activities translate to SDS reconfigurations.

  If the removed server is an SDC node, or the server to be added requires exposing storage locally, SDC reconfiguration will happen as well.
- Adding or removing storage media. Add or remove the media from the SDS with which it is associated. ScaleIO will redistribute the data accordingly and seamlessly.
- Adding or removing a node. Add or remove the SDC and SDS residing on the node. ScaleIO will redistribute the data accordingly and seamlessly.

Instant maintenance mode

Instant maintenance mode enables you to restart a server that hosts an SDS, with minimal impact on the ScaleIO system, thus bypassing the disruption and effort caused by disorderly shutdown, Protection Domain shutdown, and orderly shutdown.

Whereas ScaleIO always uses two copies of user data, invoking maintenance mode introduces an additional copy that stores all writes created during maintenance to an SDS or Fault Set (created during maintenance) in both a primary location and a new location. This copy prevents data loss if a single failure occurs.

When the SDS or Fault Set is returned from maintenance mode, only the new writes are required to be resynchronized, thus minimizing data transfer during and after the update.

Instant maintenance mode does not interrupt application IOs; it can be run on any amount of members of a Fault Set; and it can run in parallel on different Protection Domains. While an SDS is in maintenance mode, most ScaleIO operations (like adding a volume) cannot be performed in the Fault Set, Protection Domain, or Storage Pool in which the SDS and its devices reside.

To invoke maintenance mode, the following conditions are required:

- Only one Fault Unit (or standalone SDS) can be in maintenance mode at any given time.
- No other SDSs can be in degraded or failed state (force override can be used).
- There must be adequate space on other SDSs for the additional backup (force override can be used).

Note

Use of force override options when entering maintenance mode can lead to data unavailability while maintenance mode is activated.

While an SDS is in maintenance mode, it can be shut down with no danger to data.

Maintaining the virtualization layer

The following operations may be performed on volumes that are exposed by the ScaleIO virtual SAN:

- Add or remove a volume:

  Create or delete a volume in the system.
- Increase volume size:

  Add capacity to a given volume, as needed. The change in volume size occurs seamlessly without interrupting I/O.
- Map and unmap volumes to an SDC:

  This enables or disables access to a volume by an SDC, and thus by an application residing on the same node.

Management tools

You can provision, maintain, and monitor ScaleIO with the following management clients:

- Command Line Interface (CLI)

  The CLI enables you to perform the entire set of configure, maintain, and monitor activities in a ScaleIO system.
- Graphical User Interface (GUI)

  The GUI enables you to perform standard configure and maintain activities, as well as to monitor the storage system's health and performance. You can use the GUI to view the entire system, and then drill down to different elements.
- VMware plug-in (plug-in)

  The plug-in enables you to perform basic provision and maintain activities in the VMware environment. In addition, the plug-in provides a wizard to deploy ScaleIO in the VMware environment.
- OpenStack

  ScaleIO provides Cinder and Nova drivers, which enable interoperation between a ScaleIO system and an OpenStack cloud operating system.
- REST Gateway

  A REST API can be used to expose monitoring and provisioning via the REST interface. The REST server is installed as part of the ScaleIO Gateway.

  Many ScaleIO activities can be performed in more than one management tool.

  The following tool is also provided:
- Installation Manager (IM)

The IM is used for installing ScaleIO, upgrading and uninstalling components, as well as running the get-info operation. The IM is installed as part of the ScaleIO Gateway.

Configuring direct attached storage (DAS)

ScaleIO works with any free capacity—internal or direct-attached devices, either magnetic hard disk drives (HDD) or flash-based devices such as solid state drive (SSD) and PCIe cards. Although ScaleIO can work with any device topology, it is recommended to configure the raw devices as stand-alone devices.

NOTICE

SDS devices are not write-protected; their contents can be erased or formatted. Exercise caution when using disk utilities.

If the server has a RAID controller, ScaleIO prefers to use the controller's caching abilities for better performance, but is better utilized when all devices are configured as stand-alone (i.e. setting each of the devices to RAID-0 separately). For HDD devices, it is recommended to enable RAID-controller caching. As for flash devices, it depends on the device behavior.

For Windows, when using a physical disk drive, it is recommended to generate a single, unformatted partition over the entire disk.

For more information about preparing Windows devices, see the "Adding devices to SDS nodes on Windows servers" section in the *EMC ScaleIO Deployment Guide*.

Note

For HDDs: It is recommended to use RAID-controller caching when available as follows:

- READ/WRITE: if cache is battery-backed
- READ ONLY: if cache is NOT battery-backed

For flash devices (e.g. SSD): Depends on the device

PART 2

Getting Started

The chapters in this part of the guide describe how to get ScaleIO started in your environment. Chapters include:

Chapter 3, "Licensing"

Chapter 4, "User Management"

Chapter 5, "Creating and Mapping Volumes"

CHAPTER 3

Licensing

This chapter describes how to obtain and activate the electronic license for your ScaleIO software. Topics include:

- Overview 72
- Activating entitlements and installing a license file 73
- License file example 77
- Error messages 78

Overview

ScaleIO installations are enabled to be fully functional, for non-production environments.

Using ScaleIO in a production environment requires a license. The license is installed on the MDM cluster, using the SCLI `--set_license` command.

To obtain a license for production use, and to receive technical support, open a service ticket with EMC Support at https://support.emc.com.

ScaleIO licenses are purchased by physical device capacity (in TB). You can activate your licensed capacity over multiple ScaleIO systems—each system with its unique installation ID.

You download ScaleIO licenses from the EMC Software Licensing Central website, using the procedures described in "Activating entitlements and installing a license file". Then, you install the licenses on your ScaleIO system, as described in "Installing the license".

You can view current license information using the CLI or the GUI.

The following steps summarize the licensing process:

1. Purchase ScaleIO, and receive a License Authorization Code (LAC) email with a link to the licensing site.

   If you do not have the LAC email, you can search for the LAC number from the EMC Software Licensing Central website, by entering the Sales Order number and using the **Search Entitlements** option.

2. Retrieve the installation ID from your ScaleIO system.
3. Click the link in the LAC email, and use the online wizard to complete the entitlement activation process.

   a. Save the license file, and install it using the CLI.

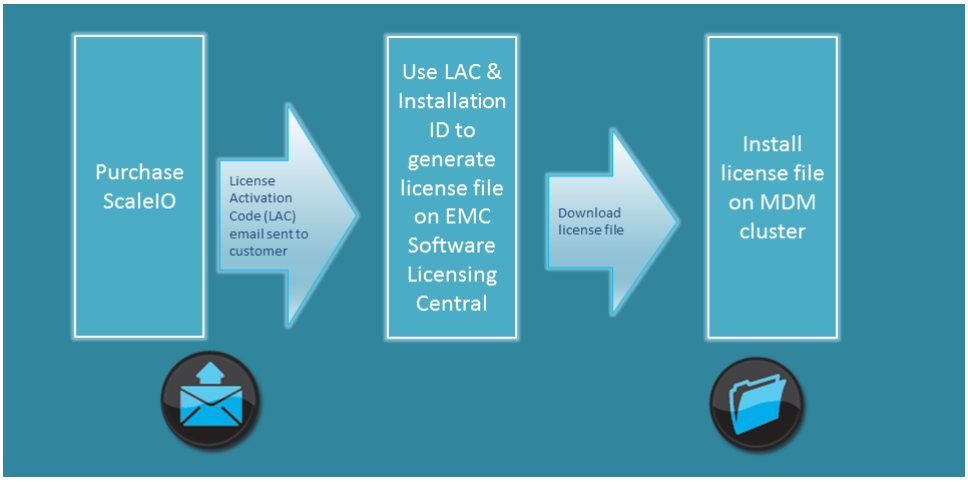

The following table describes ScaleIO eLicensing terminology.

Table 8 eLicensing terminology(continued)

Term	Description
EMC Online Support	The EMC online support portal, http://support.emc.com, contains product support information and links to the Software Licensing Central web site.
Entitlements	The EMC Software Licensing Central web site lists the entitlements (usage rights) that you have purchased, that you can activate for a specific host machine.
LAC email	Email sent to a customer who has purchased an EMC product, containing a License Authorization Code (LAC), which is needed to complete the entitlement activation process on the EMC Software Licensing Central web site.

When you purchase a license entitlement for ScaleIO, a License Authorization Code (LAC) email is sent to you, or to your purchasing department. If you cannot find the LAC email, you can use the Software Licensing Central website to find your license entitlements.

The following figure shows a sample LAC email:

Figure 11 Licensing LAC email

Dear EMC Software User,
Thank you for choosing EMC software. Your EMC Software License Authorization Code (LAC) is [redacted].
You must redeem this LAC for license keys to activate your software. Please protect your LAC like you would any other license key to prevent anyone from improperly activating your software.

Activating Your Software
1. Click here or copy and paste the following URL (https:[redacted]
LAC=[redacted]) into a web browser to activate your entitlements.
2. You will be prompted to log in. (New users should follow the new member registration steps).
3. Follow the on-screen instructions.

Downloading Your Software
1. Click here or copy and paste the following URL (https://ngtest-ci.emc.com/downloads/) into a web browser to download your software.
2. You will be prompted to log into EMC's Online Download Service Center (New users should follow the new member registration steps).
3. Enter the product name in the search field to find the software you wish to download.

License Authorization Code: RR91DYQLF9RTXZT9RMQH

Product #	Title	Quantity
456-106-154	EMC SCALEIO SOFTWARE CAPACITY=CB	250
456-106-155	EMC SCALEIO ENTERPRISE FEATURES=CB	250

If you have any questions about your sales order please contact your EMC Account Representative or your Authorized Reseller.

Activating entitlements and installing a license file

ScaleIO licenses are assigned to ScaleIO systems, each of which is identified by a unique installation ID.

You use the ScaleIO installation ID, together with your LAC, to activate the entitlement and then download the license file. Then, you install this file in your MDM cluster.

Activating an entitlement and downloading the license file

This section describes how to activate the entitlement that was purchased.

ScaleIO is procured by total capacity, but you can activate portions of this total capacity over multiple ScaleIO systems. For example, your purchase order may have been for 1000 TB. Your LAC will entitle you to activate all, or part of that. You can activate 500 TB for one ScaleIO system, and leave the rest for another activation, for the same, or a different system.

To activate the entitlement, perform the following:

Procedure

1. Identify the installation ID of your ScaleIO system:
   - Using the CLI:
     Run the following command:

     ```
     scli --query_license
     ```

     The installation ID is displayed:

     ```
     Installation ID: 0123456789abcdef
     ```

     Note

     To run CLI commands, you first need to log in. For more information see the "Logging In" section in the user documentation.

     Actual command syntax is operating-system dependent. For more information, see the *ScaleIO CLI Reference Guide*.

   - Using the GUI:
     From the top right of the main window, open the drop-down menu that appears next to the user name, and select **About**.

     The installation ID is displayed in the **About** window.
2. If you have the LAC email, skip to step 4 .
3. If you do not have your LAC email, perform the following:
   a. From the EMC support website, browse to the Software Licensing Central system:
      a. Open the EMC support website: http://support.emc.com.

         If you are a new user, create a new user account.
      b. From the **Support Tasks** list, click **Manage Licenses and Usage Intelligence**.
      c. From the software list, click **ScaleIO** . The **Powerlink Licensing** website is displayed.
      d. Click **View Entitlements**. The **Search Entitlements** screen appears.

   b. Type the Sales Order number, then click **Search Entitlements**.

      A list of entitlements is displayed.

c. Locate the entitlement to activate, and choose **Options** > **Activate**.

The **Powerlink Licensing—Search Entitlements to Activate** screen appears. Skip to step 6.

4. If you have your LAC email, perform the following:

   a. Click the link in the LAC email, and log in.

   The **Activate—Search for Products** screen appears:

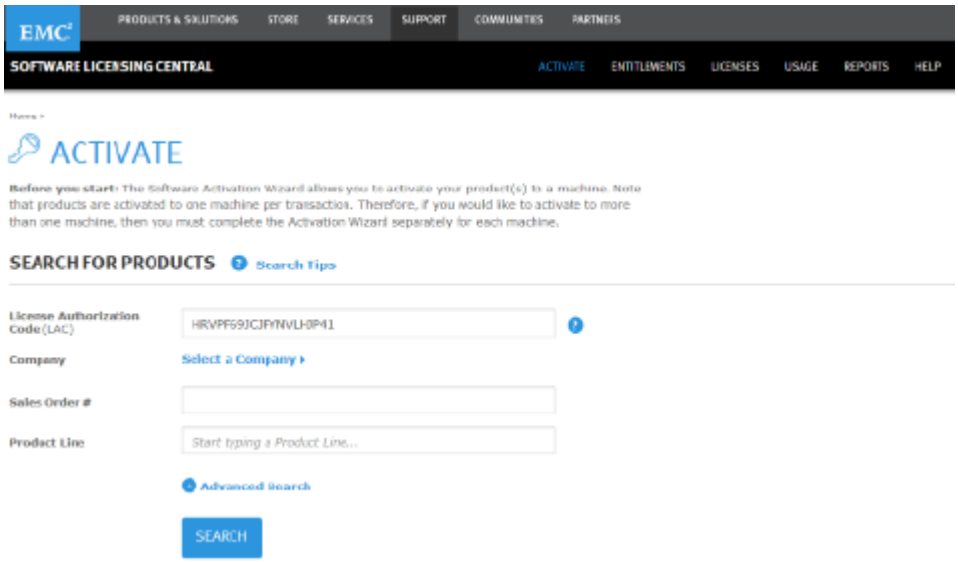

   b. Enter your LAC code, or search by Sales Order number, then click **Search**.

   The **Select Products** screen appears:

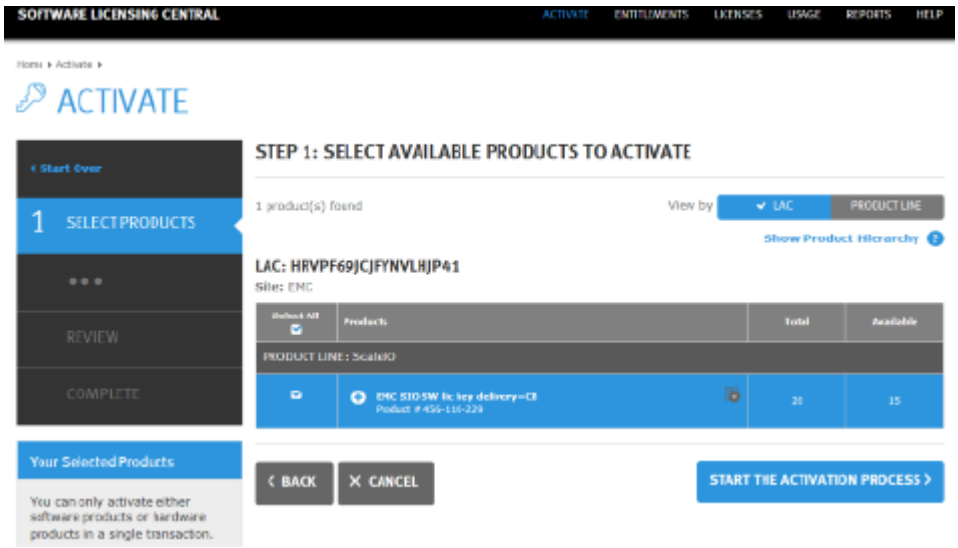

5. In the **Select Products** screen, select the product to activate, and click **Start the Activation Process**.

6. In the **Company Details** screen, confirm (or update) company information, and click **Select a Machine**.

   The **Select a Machine** screen appears:

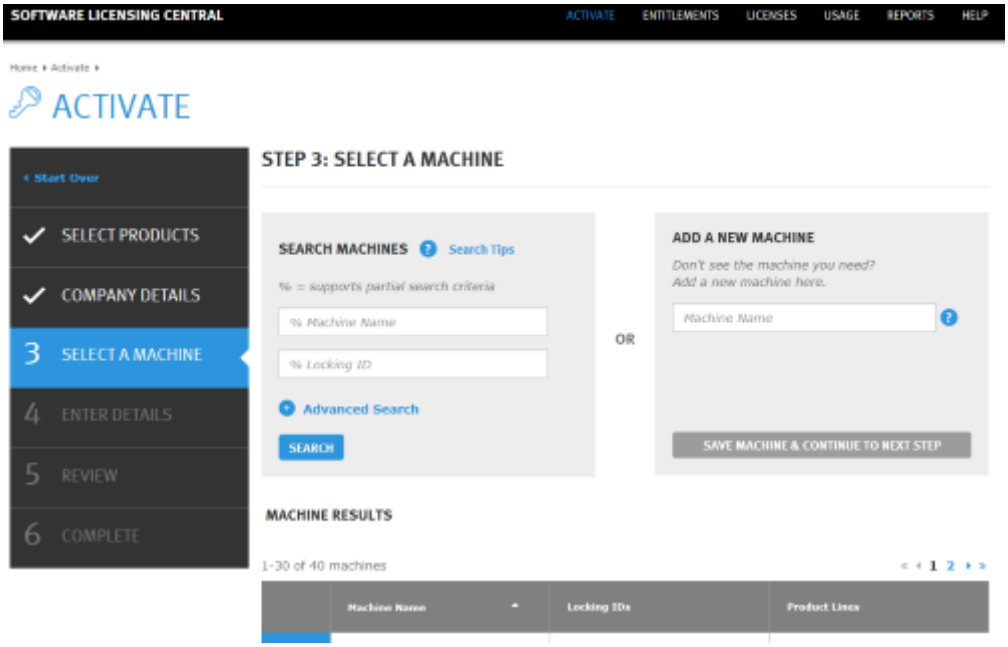

7. In the **Select a Machine** screen, select a machine on which to activate the product in one of these ways:
   - Click **Search** to locate an existing machine (one on which EMC product was previously activated).
   - Add a new machine name, then click **Save Machine & Continue**.

   In the context of the activating process, a machine is a ScaleIO system, which could comprise multiple servers.

   The **Enter Details** screen appears:

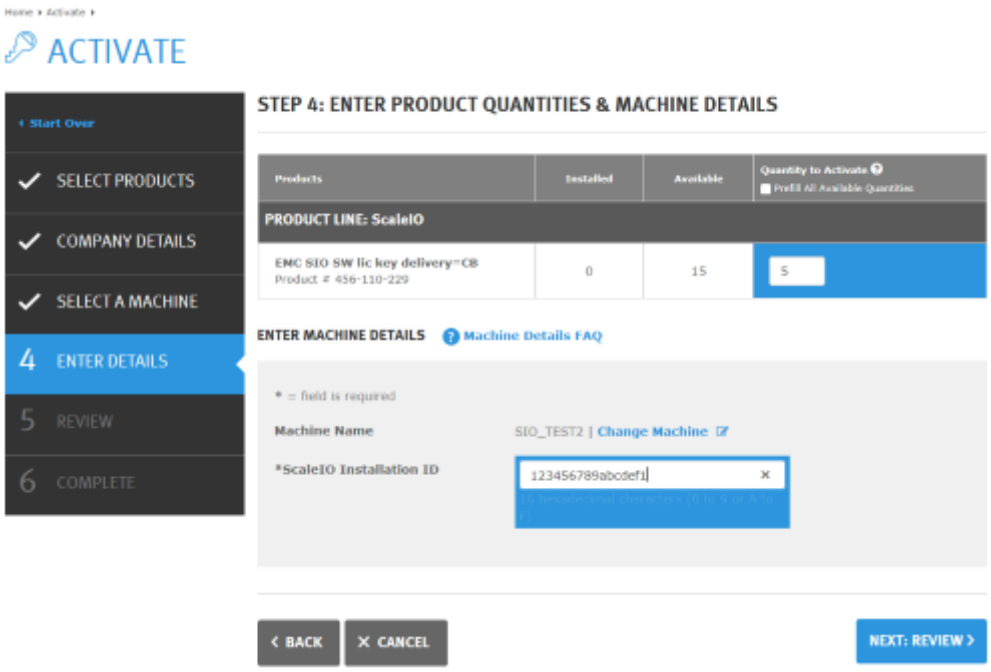

8. In the **Enter Details** screen, enter the following:
   - **Quantity** (in TB) to activate on this machine.

To allocate the available capacity over multiple machines, select less than the full amount available, and repeat the activation process on the other machine.

- ScaleIO Installation ID, from the beginning of this procedure.

9. Click **Next**.
10. In the **Review** screen, you can review your selections.

    The license key will be emailed to the user name that is logged in to the licensing system. To send it to more recipients, click **Email to more people** and enter their email addresses.
11. Click **Activate**.

Installing the license

To install the license, run the following command:

```
scli --set_license --license_file <license_file>
```

where *<license_file>* is the full path to the license file

Example: `scli --set_license --license_file /tmp/0239SH4SS89023T6.lic`

The ScaleIO license is now installed on the MDM cluster.

You can view license information using the `query_license` command and from the **About** menu in the GUI.

Actual command syntax is operating-system dependent. For more information, see the *ScaleIO CLI Reference Guide*.

License file example

The following figure illustrates a license file with a license for 200TB of capacity:

Figure 12 License file example

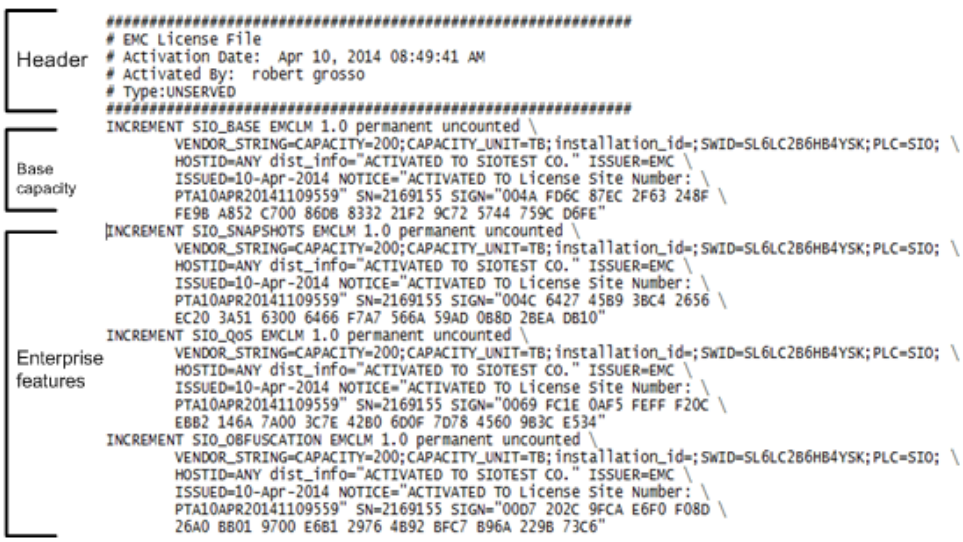

The license file includes the following sections:

- Header: General information
- General license: Shows the capacity licensed for the system, in this case 200 (TB).

Error messages

The following table lists error messages that may be generated by the system, and their troubleshooting solutions.

Table 9 Licensing error messages(continued)

Error Message	Description	Solution
The license key is invalid or does not match this version. Contact Support.	The license key is invalid.	Contact support.
The current system configuration exceeds the license entitlements.	More capacity has been installed than the license allows.	Reduce capacity, or extend the license capacity.
Operation could not be completed. The license capacity has been exceeded.	When you try to add an SDS or device, it will cause the licensed capacity to be exceeded.	Do not add the SDS or device, or extend the license capacity.
The license key is too long	The license file is larger than expected.	Check the accuracy of the license key.
The license has expired	The duration of the license has ended.	Extend the duration of the license.
The license installation ID does not match the ID of this system	When they entered the Installation ID in the ELM, it may have been incorrect	Contact support.
The license contains a mismatch of the SWID. Contact Support.	The license key is invalid.	Contact support.
The issuer of the license you are attempting to add does not match that of the product	The license key is invalid.	Contact support.
The license contains a mismatch of the capacity values for basic and advanced features. Contact Support.	The capacity licensed for basic features is not equal to the capacity licensed for advanced feature	Contact support.

CHAPTER 4

User Management

This chapter describes how to create and manage users. Topics include:

- Overview 80
- User Roles 80
- Logging in 81
- Setting the User Authentication Method 83
- Adding and modifying local users 84

Overview

ScaleIO supports local domain user authentication, and LDAP domain authentication. In addition, secure authentication is used between system internal and external components. This chapter provides the CLI commands used to create and manage ScaleIO users. The REST API can also be used to configure LDAP. For more information, see the operations for MDM clusters in "Operations" in the "REST API Reference" appendix.

- To set up local domain users, follow the instructions in this chapter.
- To set up LDAP users, see a detailed explanation in the document ScaleIO User Roles and LDAP Technical Notes. In general, the following steps must be performed:

1. Add LDAP service to the MDM.
2. Create Active Directory (AD) groups that correspond to the user roles offered by ScaleIO.
3. Set the system-wide authentication method (use with caution, because it is complex to roll-back this operation).
4. Log in again to apply the changes that you made.

User Roles

The authorization permissions of each user role are defined differently for local authentication, and for LDAP authentication. Although the role names are similar, the permissions granted to them are not. User roles defined in the LDAP domain are mutually exclusive, with no overlap (apart from the Configurator role). If you want to give an LDAP user permission to perform both monitoring and configuration roles, for example, assign that user to both Backend\Frontend Configurator and Monitor LDAP groups.

The Configurator and Super User roles do not exist at all for LDAP.

The following table describes the permissions that can be defined for local domain users and for LDAP domain users.

Table 10 Local and LDAP user roles and permissions

User role	Query		Configure parameters		Configure user credentials	
	Local	LDAP	Local	LDAP	Local	LDAP
Monitor	Yes	Yes	No	No	No	No
Configurator (this role is only applicable for local users)	Yes	Not applicable	Yes (an aggregation of both Frontend and Backend Configurator)	Not applicable	No	Not applicable
Backend Configurator	Yes	No	Yes		No	No

Table 10 Local and LDAP user roles and permissions (continued)

User role	Query		Configure parameters		Configure user credentials	
	Local	LDAP	Local	LDAP	Local	LDAP
			Backend operations only (Protection Domains, Storage Pools, Fault Sets, SDSs, Devices, other system settings)			
Frontend Configurator	Yes	No	Yes Frontend operations only (Volumes, SDCs, Snapshots)		No	No
Administrator	Yes	No	Yes	No	May configure Configurator and Monitor users	
Security Roles	No	No	No	No	May define Administrator users and control LDAP	
Super User (only one Super User is allowed per system, and it must be a local user)	Yes	Not applicable	Yes	Not applicable	Yes	Not applicable

Logging in

To access the CLI, you must first log in to the management system using a terminal application.

If the CLI and the MDM do not reside on the same server, add the `--mdm_ip` parameter to all CLI commands.

In a non-clustered environment, use the MDM IP address. In a clustered environment, use the IP addresses of the master and slave MDMs, separated by a comma. For example:

```
scli --mdm_ip 10.10.10.3,10.10.10.4 --login --username supervisor1
```

You will be prompted to enter the password.

When using LDAP, the command should also include the LDAP domain. For example:

```
scli --mdm_ip 10.10.10.3,10.10.10.4 --login --username
JohnDoe@ldap.acme.com
```

The default user created during setup is the Super User, with the *admin* user name.

login

Log the specified user into the management system. Every user must log in before performing CLI commands.

When a user is authenticated by the system, all commands will be executed with the respective role until a logout is performed, or until the session expires, by reaching one of the following timeouts:

- Maximum session length (default: 8 hours)
- Session idle time (default: 10 minutes)

Syntax

```
scli --login --username <NAME>
  [--password <PASSWORD>]
  [--ldap_authentication | --native_authentication]
  [--approve_certificate]
   --accept_banner_by_scripts_only
```

Note

Actual command syntax is operating-system dependent.

Parameters

`--username`

Username

`--password`

User password. If you do not type your password, you will be prompted to do so.

Note

In Linux, to prevent the password from being recorded in the history log, leave out the `password` flag and enter the password interactively.

`--ldap_authentication`

Log in using the LDAP authentication method. LDAP authentication parameters should be configured and LDAP authentication method should be set.

`--native_authentication`

Log in using the native authentication method (default).

`--approve_certificate`

Preemptive approval of the MDM certificate

Examples

```
scli --login --username siouser1 --password 1!2@3A
```

Note

During installation using the installation manager or the VMware plug-in, the password for the admin user is reset, and you should log in with the new password. If you installed ScaleIO manually, after logging in the first time with the default password (`admin`), you must change the password and log in again. Once that is accomplished, the admin user can create additional users.

When logging in, if a login banner has been configured and enabled in your system, you are prompted to press any key, after which the banner is displayed. To continue, you must approve the banner.

logout

Log the current user out of the system.

Syntax

```
scli --logout
```

Example

```
scli --logout
```

Setting the User Authentication Method

set_user_authentication_method

Set the user authentication method for the system.

WARNING

Use this command with caution. The operation is complex to roll back.

Note

For details about setting up LDAP, refer to the *ScaleIO User Roles and LDAP Usage Technical Notes*.

Syntax

```
scli --set_user_authentication_method (--ldap_authentication | --
native_authentication | --native_and_ldap_authentication)
  [--i_am_sure]
```

Parameters

`--ldap_authentication`

LDAP-based authentication method where users are managed on an LDAP-compliant server. Configure LDAP service parameters before switching to this authentication method.

--native_authentication

Native authentication method where users are managed locally in the system

--native_and_ldap_authentication

A hybrid authentication method. Both LDAP and Native users may log in to the system after it is set.

--i_am_sure

Skip the safety questions for command execution. (For example: “This could damage the stored data. Are you sure?”)

Example

```
scli --set_user_authentication_method --
native_and_ldap_authentication --i_am_sure
```

Adding and modifying local users

Users with administrator role can manage system users, including adding new users, deleting existing users, modifying user credentials, and resetting user passwords.

Any user with administrator role can display all of the users in the system, their roles, and their user ID, and can change the password of a user that is currently logged into the system. The following CLI commands allow you to manage the local users.

add_user

Add a user to the system. A randomly generated password for the created user is returned.

This command is available only to administrator users.

Note

When adding users, you may need to use the " or ' characters to enable the use of Bash syntax characters.

Syntax

```
scli --add_user --username <NAME> --user_role {Monitor | Configure
| BackEndConfigure | FrontEndConfigure | Security | Administrator}
```

Parameters

--username *<NAME>*

User name to add to the system

--user_role {Monitor | Configure | BackEndConfigure | FrontEndConfigure | Security | Administrator}

Role of the user: Monitor, Configurator, Backend Configurator, Frontend Configurator, Security, or Administrator. For information on user roles, see the *ScaleIO User Guide*.

Example

```
scli --add_user --username siouser2 --user_role Configure
```

delete_user

Delete the specified user from the system.

This command is available only to administrator users.

Syntax

```
scli --delete_user (--user_id <ID> | --username <NAME>)
```

Parameters

`--user_id <ID>`

ID of the user to be deleted

`--username <NAME>`

Username of the user to be deleted

Example

```
scli --delete_user --username siouser2
```

modify_user

Modify the user role of the specified user in the system.

This command is available only to administrator users.

Syntax

```
scli --modify_user (--user_id <ID> | --username <NAME>) --user_role
{Monitor | Configure | BackEndConfigure | FrontEndConfigure |
Security | Administrator}
```

Parameters

`--user_id <ID>`

User ID of the user to modify

Note

The user ID is displayed when you create the user. To find this ID at a later time, use the `query_user` command.

`--username <NAME>`

User name of the user to modify

`--user_role {Monitor | Configure | BackEndConfigure | FrontEndConfigure | Security | Administrator}`

Role of the user: Monitor, Configurator, Backend Configurator, Frontend Configurator, Security, or Administrator. For information on user roles, see the *ScaleIO User Guide*.

Example

```
scli --modify_user --username siouser3 --user_role Monitor
```

query_users

Display all the users defined in the system, with their roles and user ID.

Syntax

```
scli --query_users
```

Parameters
None.

Example

```
scli --query_users
```

query_user

Display information about the specified user.

This command is available only to administrator users.

Syntax

```
scli --query_user (--user_id <ID> | --username <NAME>)
```

Parameters

`--user_id <ID>`

User's ID number

Note

The user ID is displayed when you create the user. To find this ID at a later time, use the `query_user` command.

`--username <NAME>`

Name of the user

Example

```
scli --query_user --username sio_user
```

reset_password

Generate a new password for the specified user. The user must change the password again after logging in with the generated password.

This command is available only to administrator users.

Syntax

```
scli --reset_password (--user_id <ID> | --username <NAME>)
```

Parameters

`--user_id <ID>`

User ID of the user whose password will be reset

Note

The user ID is displayed when you create the user. To find this ID at a later time, use the `query_user` command.

`--username <NAME>`

User name of the user whose password will be reset

Example

```
scli --reset_password --username siouser3
```

set_password

Change the password of the user currently logged in to the system.

This command is available only to administrator users.

Syntax

```
scli --set_password [--old_password <OLD_PASSWORD>] [--new_password
<NEW_PASSWORD>]
```

Parameters

None.

`--old_password <OLD_PASSWORD>`

User's current password

Note

In Linux, to prevent the password from being recorded in the history log, omit the `old_password` flag and enter the password interactively.

`--new_password <NEW_PASSWORD>`

User's new password

Note

In Linux, to prevent the password from being recorded in the history log, omit the `new_password` flag and enter the password interactively.

Example

```
scli --set_password --old_password 1!2@3A --new_password P9*7&6
```

Password rules

The password must conform to the following rules:

1. Contains between six (default, configurable) and 31 characters.
2. Contains characters from at least three of the following groups: [a-z], [A-Z], [0-9], special characters (!@#$ ...)
3. The current password is not allowed.

disable_admin

Disables the default Super User.

The Super User is the default user for setting up the system, and has all the privileges of all user roles. In some cases you may need to disable the Super User in order to ensure that all users are associated with specific people.

Note

To re-enable the Super User, use the `reset_admin` command.

Syntax

```
scli --disable_admin
  [--i_am_sure]
```

Parameters

`--i_am_sure`

Skip the safety questions for command execution.

Example

```
scli --disable_admin --i_am_sure
```

Resetting the admin user password

You can reset the password of the default admin user, using a combination of a file written to the MDM, and the `reset_admin` CLI command. You must have write privileges on the MDM to complete this task.

Note

There may be multiple users defined in your system with Administrator roles. This section refers only to the default built-in admin user. You can reset the other administrator user passwords like any other user.

Command

`reset_admin`

Syntax

```
scli --reset_admin [--i_am_sure]
```

where `[--i_am_sure]` is an optional parameter which makes the operation skip the safety questions (For example: “Are you sure?”).

Actual command syntax is operating-system dependent. For more information, see “CLI basics”.

To reset the admin user password, perform the following steps:

1. Create a text file called `MDM_SERVICE_MODE` on the MDM, in the location corresponding to your operating system:
   - Windows: `C:\Program Files\emc\scaleio\MDM\logs`
   - Linux: `/opt/emc/scaleio/mdm/logs/`
2. In the body of the file, type the text **`Reset Admin`**, and save the file.
3. In the CLI, type the command `scli --reset_admin`.

   The admin user password is reset to `admin`.

CHAPTER 5

Creating and Mapping Volumes

This chapter describes how to create volumes from devices added to SDS nodes, and then to map the volumes to SDC nodes. Devices may have been added during, or after, the installation process. Topics include:

- Creating and mapping volumes overview 92
- Creating volumes 92
- Mapping a volume to an SDC 94

Creating and mapping volumes overview

You can create volumes from devices added to SDS nodes, and then map the volumes to SDC nodes. Devices may have been added during, or after, the installation process.

Creating and mapping volumes can be performed using various management tools. The following procedures describe how to do so with the GUI, and provide references to the other tools.

The creating and adding volume process described in this section is necessary, as part of the Getting Started process, before applications can access the volumes. In addition, you may create additional volumes and map them as part of the maintenance of the virtualization layer.

Creating volumes

Adding volumes

Add volumes to a system.

The adding and mapping volume process is necessary, as part of the getting started process, before applications can access the volumes. In addition, you may create additional volumes and map them as part of the maintenance of the virtualization layer.

You can create a volume when the requested capacity is available. To start allocating volumes, the system requires that there be at least three SDS nodes.

Note

For the minimum size of an SDS, see System requirements on page 26.

The created volume cannot be used until it is mapped to (at least) one SDC. For more information, see Mapping a volume to an SDC on page 94.

You can configure the caching option when creating the volumes, or you can change the Read RAM Caching feature later. If you want to enable the caching feature, ensure that the feature is also enabled in the backend of the system, for the corresponding Storage Pool and SDSs. For more information, see Changing RAM Read Cache volume settings on page 197.

Define volume names according to the following rules:

- Contains less than 32 characters
- Contains only alphanumeric and punctuation characters
- Is unique within the object type

ScaleIO objects are assigned a unique ID that can be used to identify the object in CLI commands. You can retrieve the ID via a query, or via the object's property sheet in the GUI. It is highly recommended to give each volume a meaningful name associated with its operational role.

To add one or multiple volumes, perform these steps:

Procedure

1. In the **Frontend > Volumes** view, navigate to the Storage Pool to which you want to add the volume, and select it.

2. From the **Command menu** or context-sensitive menu, select **Add Volume**.

   The **Add Volume** window is displayed.

3. If you want to create more than one volume, type the number of volumes you would like to add in the **Copies** box.

   a. If you type `1`, only one volume will be created (optional—can be left blank).

   b. If you type a number greater than 1, the characters `%i%` will be added to the **Name** box, and multiple volumes will be created, accordingly.

      The volumes will be named and numbered automatically, starting from 1. If you want the numbering to start from a different number, type it in the **Start numbering at** box, as described in Step 5. The remaining options in the window will be assigned to all the volumes created in this operation.

4. Type a name in the **Name** box.

   If you are adding multiple volumes, they will all be created with the same name, and a number will be appended instead of the characters `%i%`.

   These characters can be positioned anywhere in the name. The names that will be created are displayed in the right pane of the window, as shown in the figure later in this topic.

5. If you want the numbering to start from a specific number other than 1, type it in the **Start numbering at** box.

   This number will be the first number in the series that will be appended to the volume name. For example, if the **Name** is `Vol%i%` and the **Start numbering at** value is `100`, the name of the first volume created will be Vol100, and the second volume will be Vol101, and so on.

6. Type a number in the **Size** box, representing the volume size in GB (basic allocation granularity is 8 GB).

7. Select either **Thick** (default) or **Thin** provisioning options.

8. If obfuscation is required, select the **Use Obfuscation** check box.

9. If you want to disable the RAM Read Cache feature (disabled by default), clear the **Use RAM Read Cache** check box.

10. Click **OK**.

    The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 13 Add Volume window

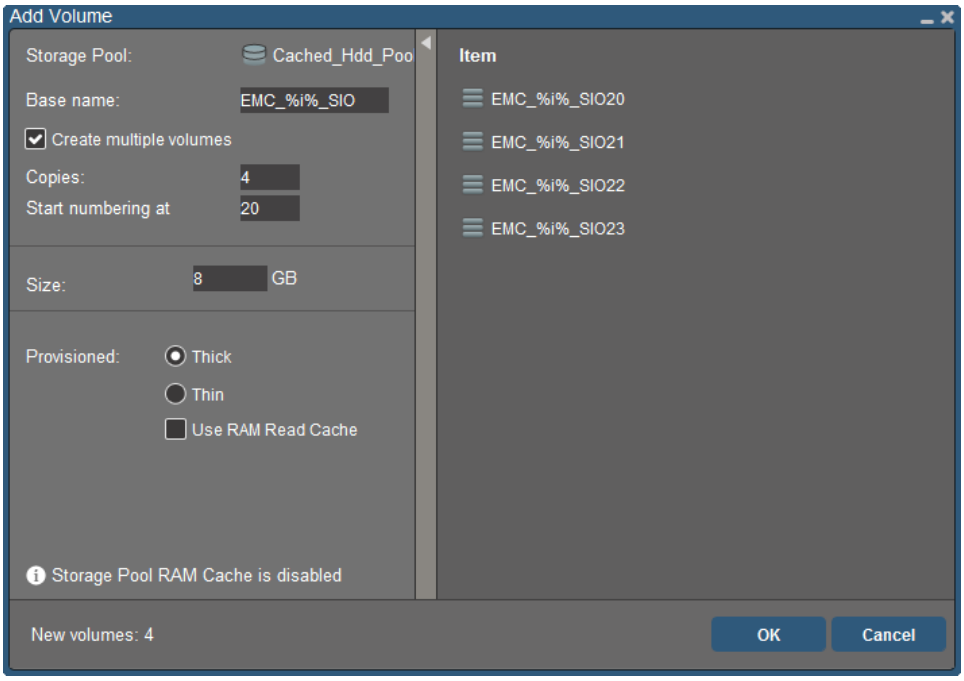

Mapping a volume to an SDC

Mapping volumes

Map one or more volumes to SDCs.

Mapping exposes the volume to the specified SDC, effectively creating a block device on the SDC.

For Linux devices, the `scini` device name can change on reboot. It is recommended to mount a mapped volume to the ScaleIO Ready Node unique ID, a persistent device name, rather than to the `scini` device name.

To identify the unique ID, run the `ls -l /dev/disk/by-id/` command. For more information, see Associating ScaleIO volumes with physical disks on page 276. You can also identify the unique ID using VMware. In the VMware management interface, device is called **EMC Fibre Channel Disk**, followed by an ID number starting with the prefix **eui**.

To map volumes, perform these steps:

Procedure

1. In the **Frontend** > **Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Map Volumes**.

   The **Map Volumes** window is displayed, showing a list of the volumes that will be mapped.
3. In the **Select Nodes** panel, select one or more SDCs to which you want to map the volumes.
   - You can use the search box to find SDCs.

- If you select an SDC that is already mapped to the volume, a green icon will appear in the mapping matrix on the right side of the window.

4. Click **Map Volumes**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

   Figure 14 Map Volumes window after mapping is complete

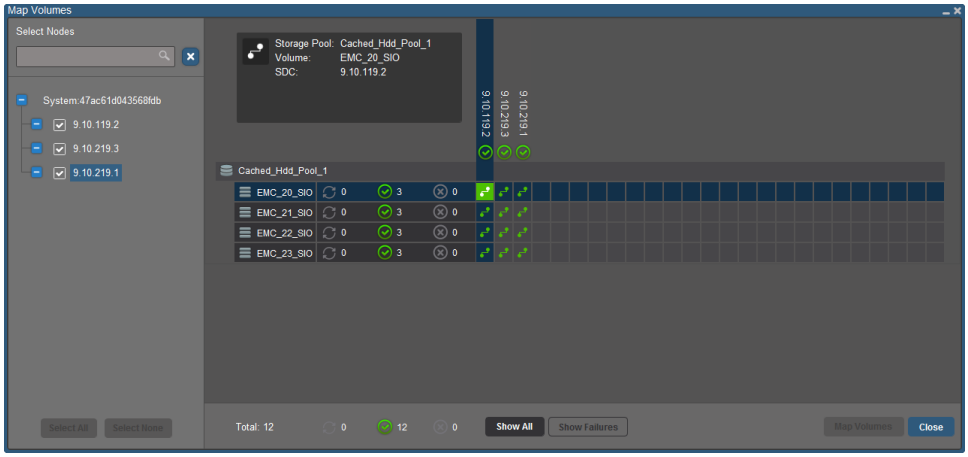

PART 3

Managing and Monitoring

By the time you have reached this part of the guide, your ScaleIO system should be up and running, all the way through mapping volumes to SDC nodes and installing the license.

The chapters in this part of the guide describe how to use the CLI to manage and monitor ScaleIO activity and components. When applicable, the use of the GUI and the VMware plug-in is also referenced. Chapters include:

Chapter 6, "Command Quick Reference"

Chapter 7, "Managing System Objects"

Chapter 8, "Security Management"

Chapter 9, "Opening the GUI and Logging In"

Chapter 10, "GUI Features"

Chapter 11, "Monitoring the System using the GUI"

Chapter 12, "Configuring the System using the GUI"

Chapter 13, "Using the VMware Plug-in"

CHAPTER 6

Command Quick Reference

Use the tables in this chapter as a quick-reference for all ScaleIO commands and actions. Each table shows the activities and monitoring options per ScaleIO component, and which management utilities support various commands.

ScaleIO CLI (scli) commands are described in the *ScaleIO CLI Reference Guide*.

ScaleIO GUI and vSphere plug-in commands are described in the *ScaleIO User Guide*.

- Device activities and monitoring 100
- MDM activities and monitoring 101
- Protection Domain activities and monitoring 103
- SDC activities and monitoring 104
- SDS activities and monitoring 106
- Storage Pool activities and monitoring 108
- User management activities and monitoring 110
- Component authentication activities and monitoring 111
- Volume activities and monitoring 112
- Cache activities and monitoring 113
- Other activities and monitoring 115
- CLI basics 116
- Syntax 118

Device activities and monitoring

Table 11 Device activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Abort device removal		`abort_remove_sds_device`	Backend > navigate to, and select the SDS > Abort Removal	
Activate a device		`activate_sds_device`	Backend > navigate to, and select the device or the SDS > Activate Device	
Add a device to an SDS		`add_sds_device`	Backend > navigate to, and select the SDS > Add device	Deployment Wizard + Advanced configuration
Clear one or more device errors		`clear_sds_device_error`	Backend > navigate to, and select the device > Clear Device Errors	Basic configuration
Modify device capacity		`modify_sds_device_capacity`	Backend > navigate to, and select the device > Set Device Capacity Limit	
Remove a device from an SDS		`remove_sds_device`	Backend > navigate to, and select the device > Remove	Basic configuration
Rename a device		`rename_device`	Backend > navigate to, and select the device > Rename	
Update original device path		`update_device_original_path`		
	Retrieve device latency meters	`query_all_device_latency_meters`	Backend > navigate to, and select the device > Device Latency table view and Property Sheet	
	Retrieve device network statistics	`query_network_latency_meters`	Backend > navigate to, and select the device > various IO table views and Property Sheet	
	Retrieve device statistics	`query_device_latency_meters`	Backend > navigate to, and select the device > Property Sheet	

Table 11 Device activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
	Retrieve device test results	`query_device_test`	Backend > navigate to, and select the device > Property Sheet	

MDM activities and monitoring

Table 12 MDM activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Add a standby MDM to a cluster		`add_standby_mdm`	Right-click the node in the Nodes panel > select Add as Standby	
Allow SDC access to the MDM		`add_sdc`		
Create an MDM cluster		`create_mdm_cluster`		
Enable\Disable restricted SDC mode		`set_restricted_sdc_mode`		
Modify the IP address of an existing MDM		`modify_management_ip`		
Prevent SDC access to the MDM		`remove_sdc`	Frontend >SDCs > Right click SDC > Remove	
Refresh a cluster's capabilities		`refresh_mdm_cluster_capabilities`		
Refresh the cluster capabilites		`refresh_mdm_cluster_capabilities`		
Remove a standby MDM from a cluster		`remove_standby_mdm`		
Rename an MDM		`rename_mdm`		

Table 12 MDM activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Replace a member of a cluster		`replace_cluster_mdm`		
Restrict remote access to MDM to read-only commands		`set_remote_read_only_limit_state`		
Set the license		`set_license`		
Set the management application ID		`set_manager_id`		
Switch the MDM Master of a cluster		`switch_mdm_ownership`		
Switch the mode of an MDM cluste		`switch_cluster_mode`	Hardware > Cluster menu > select Extend to 5 Node Cluster Mode Hardware > Cluster menu > select Decrease to 3 Node Cluster Mode	
	MDM cluster query	`query_cluster`	Backend > select System> Property Sheet	
	MDM remote access restriction state	`query_remote_read_only_limit_state`		
	Retrieve approved SDCs in restricted SDC mode	`query_all_approved_sdc`		
	Retrieve license information	`query_license`	Preferences menu > About (top right corner)	
	Retrieve restricted SDC mode state	`query_restricted_sdc_mode`		
Allow sending cluster commands during upgrade for this session		`allow_commands_during_upgrade`		

Table 12 MDM activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Set the virtual IP addresses for an MDM cluster		`modify_cluster_virtual_ips`		
Set the virtual IP interfaces for a node in the MDM cluster		`modify_virtual_ip_interfaces`		
Replace the MDM's certificate and key		`replace_mdm_security_files`		

Protection Domain activities and monitoring

Table 13 Protection Domain activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Activate a Protection Domain		`activate_protection_domain`	Backend > navigate to, and select the Protection Domain > Activate Protection Domain	
Add a Fault Set		`add_fault_set`	Backend > navigate to, and select the Protection Domain > Add Fault Set (Fault Set will only be visible in the Related Objects section of the Protection Domain's Property Sheet)	Basic configuration
Add a Protection Domain		`add_protection_domain`	Backend > select System > Add Protection Domain	Deployment Wizard + Basic configuration
Inactivate a Protection Domain		`inactivate_protection_domain`	Backend > navigate to, and select the Protection Domain > Inactivate Protection Domain	
Remove a Protection Domain		`remove_protection_domain`	Backend > navigate to, and select the	Basic configuration

Table 13 Protection Domain activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
			Protection Domain > Remove	
Remove all SDSs from a Fault Set		`clear_fault_set`		
Remove a Fault Set		`remove_fault_set`		Deployment Wizard + Basic configuration
Rename a Fault Set		`rename_fault_set`		
Rename a Protection Domain		`rename_protection_domain`	Backend > navigate to, and select the Protection Domain > Rename	
	Single Fault Set query	`query_fault_set`		
	Fault Sets query	`query_all_fault_sets`		
	Protection Domain query	`query_protection_domain`	• Dashboard > Cluster > Protection Domains pane • Backend > navigate to, and select the Protection Domain > Property Sheet	Used for viewing

SDC activities and monitoring

Table 14 SDC activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in	Notes
	All SDCs query	`query_all_sdc`	Frontend > SDCs > select the SDC and display the Property Sheet	Used for viewing	
Remove SDC		`remove_sdc`	Frontend > SDCs > right-click the SDC and select Remove		
Rename an SDC		`rename_sdc`	Frontend > SDCs > right-click the SDC and select Rename		

Table 14 SDC activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in	Notes
Add an MDM definition to an SDC (usually to add the SDC to an additional SIO system)		`drv_cfg --add_mdm`			The drv_cfg command line is a local CLI utility that affects only the client on which the SDC is running (unlike the ScaleIO CLI which may affect the entire system). Refer to the user documentation for the correct usage.
Load an MDM configuration file to an SDC		`drv_cfg --load_cfg_file`			
Modify an existing MDM IP's address configured on an SDC		`drv_cfg --mod_mdm_ip`			
Modify MDM IP addresses or GUID configured on an ESX-based SDC		`esxcli system module parameters list -m scini`			The esxcli command line is a local CLI utility used on an ESX server that affects only the client on which the SDC is running (unlike the ScaleIO CLI, which may affect the entire system). Refer to the utility documentation for the correct usage.
	Query GUID and MDM IP addresses on an ESX-based SDC	`esxcli system module parameters list -m scini`			
	Query SDC state on an ESX-based SDC	`esxcli system module list |grep scini`			
Abort SDC removal		`abort_remove_sdc`			
	SDC query	`query_sdc`			
	Query all active tgt objects	`drv_cfg --query_tgts`			The drv_cfg command line is a local CLI utility that affects only the client on which the SDC is running (unlike the ScaleIO CLI which may affect the entire system). Refer to the user documentation for the correct usage.
	Query GUIDs	`drv_cfg --query_guid`			
	Query to determine to which MDM an SDC is connected	`drv_cfg --query_mdms`			

SDS activities and monitoring

Table 15 SDS activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in	Notes
Add an IP address to an SDS		`add_sds_ip`	Backend > navigate to, and select SDS > Configure IP Addresses		
Add SDS nodes		`add_sds`	Backend > navigate to, and select a Protection Domain > Add SDS	Deployment Wizard	
Change the SDS port used by SIO		`modify_sds_port`	Backend > navigate to, and select SDS > Configure IP Addresses		
Enter Maintenance Mode		`enter_maintenance_mode`	Backend > navigate to, and select SDS or Fault Set > Enter Maintenance Mode		
Exit Maintenance Mode		`exit_maintenance_mode`	Backend > navigate to, and select SDS or Fault Set > Exit Maintenance Mode		
Modify an SDS IP role		`modify_sds_ip_role`	Backend > navigate to, and select SDS > Configure IP Addresses		
Name\Rename an SDS		`rename_sds`	Backend > navigate to, and select SDS > Rename		
Remove an SDS node		`remove_sds`	Backend > navigate to, and select SDS > Remove		
Remove an IP address from an SDS		`remove_sds_ip`	Backend > navigate to, and select SDS > Configure IP Addresses		
Set SDS rebalance and rebuilt limits		`set_sds_network_limits`	Backend > navigate to, and select Protection Domain > Set Network Throttling (configures all SDSs		

Table 15 SDS activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in	Notes
			in the specified Protection Domain)		
Test SDS device		`start_device_test`			
Test the SDS network		`start_sds_network_test`			
	All SDSs query	`query_all_sds`	Dashboard > SDSs tile Backend > Filter by SDSs and toggle between various table views		
	Query volumes	`drv_cfg --query_vols`	Dashboard > Volumes tile Backend > filter by Volumes > toggle between various table views		The drv_cfg command line is a local CLI utility that affects only the client on which the SDC is running (unlike the ScaleIO CLI which may affect the entire system). Refer to the user documentation for the correct usage.
	Retrieve network test results	`query_sds_network_test_results`			
	Single SDS query	`query_sds`	Backend > Filter for the SDS > select SDS in table and either open its Property Sheet or toggle between various table views		
	Query SDS connectivity status	`query_sds_connectivity_status`	Dashboard > SDSs tile Backend > Filter by SDSs and toggle between various table views		
Abort the removal of an SDS		`abort_remove_sds`	Backend > navigate to, and right click the SDS > Abort Removal		

Table 15 SDS activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in	Notes
	Query SDS device	`query_sds_device_info`	Backend > Navigate to and select device > display Property Sheet		
Set the properties of the dirty-region-logging		`set_drl_properties`	Backend > right-click SDS > Select DRL Mode (Volatile,Hardened)		
Turns the device LED on or off		`set_sds_device_led`			

Storage Pool activities and monitoring

Table 16 Storage Pool activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Add a Storage Pool		`add_storage_pool`	Backend > navigate to, and select a Protection Domain > Add Storage Pool	Deployment Wizard + Basic configuration
Modify spare capacity policy		`modify_spare_policy`	Backend > navigate to, and select the Storage Pool > Configure Spare Percentage Policy	
Modify Storage Pool's zero padding policy		`modify_zero_padding_policy`	Backend > navigate to, and select the Storage Pool > Modify Zero Padding	
Remove a Storage Pool		`remove_storage_pool`	Backend > navigate to, and select the Storage Pool > Remove	Basic configuration
Rename a Storage Pool		`rename_storage_pool`	Backend > navigate to, and select the Storage Pool > Rename	
Set the rebuild policy		`set_rebuild_policy`	Backend > navigate to, and select the Storage Pool > Set I/O Priority	
Set the rebuild mode		`set_rebuild_mode`	Backend > navigate to, and select the Storage Pool > Enable/Disable Rebuild/Rebalance	

Table 16 Storage Pool activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Set the rebalance policy		`set_rebalance_policy`	Backend > navigate to, and select the Storage Pool > Set I/O Priority	
Set the rebalance mode		`set_rebalance_mode`	Backend > navigate to, and select the Storage Pool > Enable/Disable Rebuild/Rebalance	
Set parallelism limit (concurrent rebuild and rebalance activities)		`set_rebuild_rebalance_parallelism`		
Enable background device scanner		`enable_background_device_scanner`	Backend > navigate to, and select the Storage Pool > Set Background Device Scanner Mode	
Disable background device scanner		`disable_background_device_scanner`	Backend > navigate to, and select the Storage Pool > Set Background Device Scanner Mode	
Reset background device scanner error counters		`reset_scanner_error_counters`	Backend > navigate to, and select the Storage Pool > Reset Background Device Scanner Counters	
Set use of in-flight checksum validation		`set_checksum_mode`	Backend > navigate to, and select a Storage Pool > Configure Use Checksum	
	Storage Pool query	`query_storage_pool`	• Dashboard > navigate to Protection Domain node > Storage Pools tile • Backend > Filter for Storage Pool > select the Storage Pool in table and either open its Property Sheet or toggle various table views	

User management activities and monitoring

Table 17 User management activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Add a user		`add_user`		
Add LDAP service		`add_ldap_service`		
Assign LDAP groups to roles		`assign_ldap_groups_to_roles`		
Remove LDAP group assignments		`remove_ldap_group_from_role_assignment`		
Change a user's password		`set_password`		
Delete a user		`delete_user`		
Disable the Super User role		`disable_admin`		
Login		`login`	Start the application; use the Login window	
Logout		`logout`	Preferences menu > Logout (top right corner)	
Modify user credentials		`modify_user`		
Remove LDAP service		`remove_ldap_service`		
Rename LDAP service		`rename_ldap_service`		
Reset a user's password		`reset_password`		
Reset an admin user (special conditions required)		`reset_admin`		
Set user authentication method		`set_user_authentication_method`		
	Display information about a user	`query_user`		
	Display information about user authentication	`query_user_authentication_properties`		

Table 17 User management activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
	Display system users and their roles	`query_users`		

Component authentication activities and monitoring

Table 18 Component authentication activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Generating an MDM Certificate		`generate_mdm_certificate`		
Adding a certificate to the CLI's Trusted Certificate list		`add_certificate`		
Adding certificates to the GUI's Trusted Certificate list			Preferences menu (top right corner) > System Settings > Renew Certificates	
Setting component authentication properties		`set_component_authentication_properties`		
Generating an SDS certificate		`generate_certificate`		
Approving all MDM certificates		`approve_all_mdm_certificates`		
Generating an MDM CSR file		`generate_mdm_csr_file`		
Setting management client communication		`set_management_client_communication`		
	Showing an SDS certificate	`set_management_client_communication`		
	Showing information about the SSL version used by the MDM cluster		System > Property Sheet > MDM Cluster	
	Show pending cluster certificates		Alerts view	

Table 18 Component authentication activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
			Backend view > State Summary	
Setting preemptive acceptance of the login banner		`set_cli_login_banner_preemptive_acceptance`		

Volume activities and monitoring

Table 19 Volume activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Create a volume		`add_volume`	Frontend > Volumes > Add Volume	Advanced configuration
Detect new volumes		`drv_cfg --rescan`[a]		
Increase volume capacity		`modify_volume_capacity`	Frontend > Volumes > navigate to the volumes > Increase Volumes' Size	
Map a volume to an SDC		`map_volume_to_sdc`	Frontend > Volumes > Map Volume	Advanced configuration
Remove a volume		`remove_volume`	Frontend > Volumes > navigate to the volumes > Remove	Advanced configuration
Remove consistency group snapshots		`remove_consistency_group_snapshots`	Frontend > Snapshots > navigate to the volume > Remove Descendants only or Remove with Descendants or Remove Consistency Group	
Rename a volume		`rename_volume`	Frontend > Volumes > right-click the volume > Rename	
Set volume bandwidth limits		`set_sdc_volume_limits`	Frontend > Volumes > navigate to the volumes > Set Volume Limits	
Take a snapshot of a volume		`snapshot_volume`	Frontend > Volumes > Snapshot Volume	

Table 19 Volume activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Unmap a volume		`unmap_volume_from_sdc`	Frontend > Volumes > Unmap Volume	Advanced configuration
	All volumes query	`query_all_volumes`	Dashboard > System Overview> Volumes tile	
	Single volume query	`query_volume`	Frontend > Volumes > select the volume and open the Property Sheet	
	Volume tree query	`query_volume_tree`		
	Retrieve volume bandwidth limits	`query_sdc_volume_limits`	Frontend > SDCs > select an SDC and open the Property Sheet > Mapped Volumes	

a. The drv_cfg command line is a local CLI utility that affects only the client on which the SDC is running (unlike the ScaleIO CLI which may affect the entire system). Refer to the user documentation for the correct usage.

Cache activities and monitoring

Table 20 Cache activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Set volume RAM Read Cache (Note that RAM Read Cache is disabled for Version 2.0.x and higher)		`set_volume_rmcache_usage`		
Control the RAM cache usage in a Storage Pool		`set_rmcache_usage`	Backend > navigate to, and select Storage Pool > Configure RAM Read Cache	
Disable RAM cache of the SDS		`disable_sds_rmache`	Backend > navigate to, and select SDS > Configure RAM Read Cache	
Enable RAM cache of the SDS		`enable_sds_rmcache`	Backend > navigate to, and select SDS > Configure RAM Read Cache	
Set the RAM cache size of the SDS		`set_sds_rmcache_size`	Backend > navigate to, and select SDS >	

Table 20 Cache activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
			Configure RAM Read Cache	
Set the RAM cache write handling mode of the Storage Pool		`set_rmcache_write_handling_mode`	Backend > navigate to, and select the Storage Pool > Configure RAM Read Cache	
Control the Flash Read Cache usage in a Storage Pool		`set_rfcache_usage`	Backend > navigate to, and select the Storage Pool > Set Flash Cache Policy	
Enable Read Flash Cache of an SDS		`enable_sds_rfcache`	Backend > navigate to, and select SDS > Set Flash Cache Policy	
Disable Read Flash Cache of an SDS		`disable_sds_rfcache`	Backend > navigate to, and select SDS > Set Flash Cache Policy	
Configure a Read Flash Cache device		`set_rfcache_parameters`		
Add a Read Flash Cache device to an SDS		`add_sds_rfcache_device`	Backend > navigate to, and select SDS > Add Read Flash Cache Device	
Remove a Read Flash Cache device from an SDS		`remove_sds_rfcache_device`	Backend > navigate to, and select SDS > Remove Read Flash Cache Device	
Clear Read Flash Cache device errors		`clear_sds_rfcache_error`	Backend > right-click the device > Clear Device Errors	
Restore the path that was originally configured for the Read Flash Cache device		`update_sds_rfcache_device_original_path`		
Rename the Read Flash Cache device		`rename_sds_rfcache_device`	Backend > right-click the device > Rename	

Other activities and monitoring

Table 21 Other activities and monitoring

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
Configure Oscillating Failure counter properties		`set_oscillating_failure_counter_parameters`	Backend > System | Protection Domain | Storage Pool > Set Oscillating Failure Properties	
Configure the facility field for syslog events		`set_syslog_facility`		Deployment Wizard
Get help on a CLI command		`help`		
Rename the system		`rename_system`	Backend > select System table row > Rename	
Reset Oscillating Failure counters		`reset_oscillating_failure_counters`	Backend > System | Protection Domain | Storage Pool > Reset Oscillating Failure Counters	
Set capacity threshold alert		`set_capacity_alerts_threshold`	Backend > select System > Set Capacity Alert Thresholds	
Set use of performance profiles		`set_performance_parameters`	Backend or Frontend > select object > Set Performance Profile for...	
Show events		`showevents.py`		
Start sending events to a remote syslog server		`start_remote_syslog`		
Stop sending events to a remote syslog server		`stop_remote_syslog`		
Start an upgrade		`start_upgrade`		
Stop an upgrade		`abort_upgrade`		
Finalize an upgrade		`finalize_upgrade`		
	Upgrade query	`query_upgrade`		
	General query	`query_all`	Dashboard	

Table 21 Other activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in
			Backend > Overview or Configuration table views > Property Sheets for more details	
	General property query	`query_properties`	Use the Property Sheet	
	Oscillating Failure counters	`query_oscillating_failure_counter_parameters`	Network: Backend > System > Download Network Failure Counters SDCs: Frontend > SDCs > Property Sheet > Oscillating Failure Counters SDSs: Backend > SDSs > Property Sheet > Oscillating Failure Counters	
	Performance parameter usage query	`query_performance_parameters`	Backend > System > Property Sheet > Performance	
	Remote syslog query	`query_remote_syslog`		
	Various system limits	`query_system_limits`	Backend > System > Property Sheet	
Sets or remove the login banner		`set_login_banner`		
Enable or disable preemptive acceptance of login banner		`set_cli_login_banner_preemptive_acceptance`		
	The certificate of the given SDSs	`show_certificate`		

CLI basics

The ScaleIO CLI, SCLI, enables you to perform all provision, maintain, and monitor activities.

The CLI is installed as part of the MDM component and can be found in the following path:

- Linux and VMware: `scli`

- Xen: `siocli`
- Windows: `C:\Program Files\emc\scaleio\MDM\bin`

All CLI commands use the following format:

- Linux, VMware, and Windows:

```
scli [--mdm_ip <IP>] <command>
```

- Xen:

```
siocli [--mdm_ip <IP>] <command>
```

Description: Execute a CLI command.

Parameter	Description
--mdm_ip <IP>	One, or more IP addresses of the servers running the Master MDM and Slave. In a non-clustered environment, use the MDM IP address. If the CLI does not reside on the MDM, the `--mdm_ip` parameter must be added to every CLI command.
--approve_certificate	Preemptive approval of the MDM's certificate
<command>	Command to be executed

```
scli --mdm_ip 10.10.10.3,10.10.10.4 --query_all
```

The `mdm_ip` indicates the MDM that receives and is to execute the command. If the command is run from the Master MDM, this switch may be omitted.

To avoid using the `--mdm_ip` parameter in every command, or avoid having to install the CLI on other servers, use SSH or RDM to log in to the shell running on the management server.

You cannot execute SCLI commands on the Slave MDM. However, you can send a command from the Slave MDM by adding the IP address of the Master MDM to the command, using the `--mdm_ip` parameter.

Note

- The order of the parameters and command is insignificant.
- SCLI commands are lowercase and case-sensitive.
- All parameters are preceded by `--`

Before using most SCLI commands, you must log in, as described in the *CLI Reference Guide*.

Using SCLI in non-secure mode

If ScaleIO is running in non-secure mode, you must make the following change to enable running commands.

Note

For more information on how to set up secure or non-secure mode, see the *ScaleIO Ready Node AMS User Guide* .

On every MDM server, disable secure communication:

- Windows

  In the SCLI `conf.txt` file, add `cli_use_secure_communication=0`
- Linux

  Run this command: `echo cli_use_secure_communication=0 >> ~/.scli/conf.txt`

Note

After renewing certificates, changing from secure mode to non-secure mode is not supported.

Syntax

All names of objects in the system will be capitalized, for example, Protection Domain. In the case where the name is in fact initials it will be in uppercase, for example, MDM.

The actual CLI command format uses the following format:

- Message - Required
- <> - Argument
- () – Required element
- [] – Optional element
- | - Select from options A|B|C

  Each command entry uses the above syntax and looks like the following example:

  Usage:

```
scli --cmd_example --r1 (--r2 | --r3 <V1>) [o1 <V2>|o2]
[Options]
```

Description: a description of what `cmd_example` does

Parameter	Description
--r1	r1 description
--r2	r2 description
--r3 <V1>	r3 description with possible V1 input values
--o1 <V2>	o1 description with possible V2 input values
--o2	o2 description
Options: CHOOSE SEVERAL	
--so1 <V3>	so1 description with possible V3 input values

Parameter	Description
--so2	so2 description

The interpretation is as follows:

- The text `scli --cmd_example r1` is mandatory.
- `(--r2 | --r3 <V1>)` indicates that you must choose one of the options separated by "|". Selecting an option is REQUIRED, indicated by "( )".
- `[o1 <V2>|o2]` indicates that you may choose one of the options separated by "|". Selecting an option is OPTIONAL, indicated by "[ ]".
- `[Options]` indicates that you may choose one of the options that will be described in the table under Options. It is OPTIONAL, indicated by "[ ]".

Getting help with the CLI

The CLI supports auto-completion. To complete a command or parameters, press the TAB key while typing CLI commands.

Note

In Windows, ScaleIO does not support auto-completion.

Alternately, you can run the `help` command.

Command

`help`

Syntax

```
scli --help [Options]
```

Description/Notes

Use this command to view CLI help.

Parameters

Parameter	Description
Options: CHOOSE ONE	
`mdm`	Help on MDM commands
`sds`	Help on SDS commands
`vol`	Help on volume-related commands
`sdc`	Help on SDC commands
`general`	Help on general commands
`all`	Help on all commands

Parameter	Description
`user`	Help on commands related to user management
`<blank>`	List options

Example

```
scli --help --mdm
```

CHAPTER 7

Managing System Objects

This chapter describes how to manage and configure ScaleIO system objects. Topics include:

- Extend an existing ScaleIO system 122
- Managing the MDM cluster 122
- Managing the SDSs and cache 131
- SDC operations 134
- Managing ESX servers 142

Extend an existing ScaleIO system

Options for adding nodes to an existing system.

You can add nodes to an existing system, as well as extend the MDM cluster from a 3-node to a 5-node cluster. Depending on your system, you can use the Installation Manager (for physical servers) or the vSphere plug-in (for ESXi servers).

These topics are described in the sections of the *ScaleIO Deployment Guide*:

- "Extending an existing ScaleIO system"
- "Extending the MDM cluster from 3 to 5-node"

Managing the MDM cluster

You can replace or update the server IP address of an MDM that is currently a member of the MDM cluster. Various management tools are available to configure virtual IP addresses for the MDMs.

Replacing, or updating an IP address on a member of the MDM cluster

This section describes how to replace an MDM server that is a member of the MDM cluster. This could be necessitated by a need to replace a faulty server or a need to change the server IP address. For purposes of this section, we will refer to the server that needs to be replaced as the current server.

There are two ways to accomplish this task, determined by whether you are able to add a new server to the cluster, or whether you do not have an extra server to add. If you have an extra server to replace the current server, then there is no need to change the cluster mode (3-node or 5-node). However, if you do not have an additional server, you will need to reduce the cluster mode, from 5-node to 3-node, or from 3-node to single node.

Note

It is not recommended to use single mode in production systems, except in temporary situations.

Regardless of the circumstances, the following rules are true:

- To remove a cluster member, you first make it a standby, then remove the standby.

  To add a member to a cluster, you first make it a standby, then add the standby to the cluster. In other words, you cannot move a server from being a cluster member to being entirely external, in either direction, without being a standby first.
- The cluster must always have 5, 3, or 1 members, never any other amount.

  For a further understanding of this subject, see The MDM cluster on page 35.

  Proceed to the section that describes your environment:

  - Replacing a cluster member by adding a new member to the cluster on page 123
  - Replacing a cluster member without adding a new server to the cluster on page 124

Replacing a cluster member by adding a new member to the cluster

This section describes how to replace a member of the cluster, by adding a new member to the cluster to take its place.

Before you begin, perform the following:

- Assign the necessary IP addresses to the replacement server.
- Install the MDM package on the server.

  In this example, we are replacing the server whose IP address is 10.103.110.179, currently a member of a 5-node MDM cluster, with a server a server whose IP address is 10.103.110.57, which is currently external to any ScaleIO system. This process can be used to replace any role in the MDM cluster.

Procedure

1. Ensure that the current server (179) is not the Master MDM, by running the following command:

   ```
   scli --query_cluster
   ```

   Output, similar to the following, is displayed:

   ```
   # scli --query_cluster
   Cluster:
   Mode: 5_node, State: Normal, Active: 5/5, Replicas: 3/3
   Master MDM:
   Name: mdm17, ID: 0x5d07497754427fd0
   IPs: 10.103.110.17, 192.168.1.17, Management IPs:
   10.103.110.17, Port: 9011
   Version: 2.0.972
   Slave MDMs:
   Name: mdm19, ID: 0x26ee566356362451
   IPs: 10.103.110.19, 192.168.1.19, Management IPs:
   10.103.110.19, Port: 9011
   Status: Normal, Version: 2.0.972
   Name: mdm18, ID: 0x5843c4d16d8f1082
   IPs: 10.103.110.18, 192.168.1.18, Management IPs:
   10.103.110.18, Port: 9011
   Status: Normal, Version: 2.0.972
   Tie-Breakers:
   Name: mdm179, ID: 0x7380b70e2f73d346
   IPs: 10.103.110.179, 192.168.1.179, Port: 9011
   Status: Normal, Version: 2.0.972
   Name: mdm20, ID: 0x6dfe1c5f4062b5b3
   IPs: 192.168.1.20, 10.103.110.20, Port: 9011
   Status: Normal, Version: 2.0.972
   ```

   In this case, server 179 is a Tie Breaker.

2. If the current server is the Master MDM, change its state using the `switch_mdm_ownership` command, as described in the *ScaleIOCLI Reference Guide*.

3. Make the replacement MDM server a standby MDM, and assign it a name (*mdm57*, in our example) by running the following command, on the Master MDM:

```
scli --add_standby_mdm --mdm_role tb --new_mdm_ip
10.103.110.57,192.168.1.57 --new_mdm_management_ip
10.103.110.57 --new_mdm_name mdm57
```

4. You can see the result of the command by running the following command:

```
scli --query_cluster
```

Output, similar to the following, is displayed:

```
# scli --query_cluster
Cluster:
Mode: 5_node, State: Normal, Active: 5/5, Replicas: 3/3
...
Tie-Breakers:
Name: mdm179, ID: 0x7380b70e2f73d346
IPs: 10.103.110.179, 192.168.1.179, Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.103.110.20, Port: 9011
Status: Normal, Version: 2.0.972
Standby MDMs:
Name: mdm57, ID: 0x073e4c8b1d20d124, Tie Breaker
IPs: 10.103.110.57, 192.168.1.57, Port: 9011
```

mdm57 has been added as a standby MDM. Once it is a standby MDM, it can be added to the cluster.

5. Replace the current *mdm179* with the standby *mdm57* by running the following command:

```
scli --replace_cluster_mdm --remove_tb_name mdm179
--add_tb_name mdm57
```

The following output is displayed:

```
Successfully replaced the cluster MDM
```

The current server has been replaced.

Replacing a cluster member without adding a new server to the cluster

This section describes how to replace a member of the cluster by removing it from the cluster, and then adding it back to the cluster. This procedure requires reducing the amount of nodes in the MDM cluster.

In this example, we are removing the current server whose IP address is 10.103.110.179, currently a Tie Breaker member of a 5-node MDM cluster. Because we must retain a majority in the MDM cluster, we must also remove one of the Slave MDMs in the cluster, in this case the MDM whose IP address is 10.103.110.19. This process can be used to replace any role in the MDM cluster.

Procedure

1. Ensure that the current server (179) is not the Master MDM:

```
scli --query_cluster
```

Output, similar to the following, is displayed:

```
# scli --query_cluster
Cluster:
Mode: 5_node, State: Normal, Active: 5/5, Replicas: 3/3
Master MDM:
Name: mdm17, ID: 0x5d07497754427fd0
IPs: 10.103.110.17, 192.168.1.17, Management IPs:
10.103.110.17, Port: 9011
Version: 2.0.972
Slave MDMs:
Name: mdm19, ID: 0x26ee566356362451
IPs: 10.103.110.19, 192.168.1.19, Management IPs:
10.103.110.19, Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm18, ID: 0x5843c4d16d8f1082
IPs: 10.103.110.18, 192.168.1.18, Management IPs:
10.103.110.18, Port: 9011
Status: Normal, Version: 2.0.972
Tie-Breakers:
Name: mdm179, ID: 0x7380b70e2f73d346
IPs: 10.103.110.179, 192.168.1.179, Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.103.110.20, Port: 9011
Status: Normal, Version: 2.0.972
```

In this case, server 179 is a Tie Breaker.

2. If the current server is the Master MDM, change its state using the `switch_mdm_ownership` command, as described in the *ScaleIO CLI Reference Guide*.
3. Switch to a 3-node cluster:

```
scli --switch_cluster_mode --cluster_mode 3_node
--remove_tb_name mdm179 --remove_slave_mdm_name mdm19
```

The following output is displayed:

```
Successfully switched the cluster mode.
```

4. To view the result of the command, run:

```
scli --query_cluster
```

Output similar to the following is displayed:

```
# scli --query_cluster
Cluster:
Mode: 3_node, State: Normal, Active: 3/3, Replicas: 2/2
...
```

```
Slave MDMs:
Name: mdm18, ID: 0x5843c4d16d8f1082
IPs: 10.103.110.18, 192.168.1.18, Management IPs:
10.103.110.18, Port: 9011
Status: Normal, Version: 2.0.972
Tie-Breakers:
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.103.110.20, Port: 9011
Status: Normal, Version: 2.0.972
Standby MDMs:
Name: mdm19, ID: 0x26ee566356362451, Manager
IPs: 10.103.110.19, 192.168.1.19, Management IPs:
10.103.110.19, Port: 9011
Name: mdm179, ID: 0x7380b70e2f73d346, Tie Breaker
IPs: 10.103.110.179, 192.168.1.179, Port: 9011
```

The cluster has been changed to 3-node mode, as a Slave MDM (*mdm19*) and a TB MDM (*tb179*) have been removed and are now standby MDMs.

Now that the current server is a standby MDM, it can removed from the ScaleIO system.

5. Remove the current server from the ScaleIO system:

```
scli --remove_standby_mdm --remove_mdm_name mdm179
```

The following output is displayed:

```
Successfully removed the standby MDM.
```

6. To view the result of the command, run:

```
scli --query_cluster
```

Output similar to the following is displayed:

```
Cluster:
Mode: 3_node, State: Normal, Active: 3/3, Replicas: 2/2
...
Tie-Breakers:
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.103.110.20, Port: 9011
Status: Normal, Version: 2.0.972
Standby MDMs:
Name: mdm19, ID: 0x26ee566356362451, Manager
IPs: 10.103.110.19, 192.168.1.19, Management IPs:
10.103.110.19, Port: 9011
```

The current server is no longer a standby MDM.

7. Reassign IP addresses to the current server, as required.

In our case, we will assign the following IP address to the current server: 10.103.110.57.

8. Add the current server (57) back to the system as a standby MDM, and assign it the name `mdm57`:

```
scli --add_standby_mdm --mdm_role tb --new_mdm_ip
10.103.110.57,192.168.1.57 --new_mdm_management_ip
10.103.110.57 --new_mdm_name mdm57
```

Output similar to the following is displayed:

```
Successfully added a standby MDM. Object ID 13c925450656db74
```

9. To view the result of the command, run:

```
scli --query_cluster
```

Output similar to the following is displayed:

```
Cluster:
Mode: 3_node, State: Normal, Active: 3/3, Replicas: 2/2
...
Tie-Breakers:
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.103.110.20, Port: 9011
Status: Normal, Version: 2.0.972
Standby MDMs:
Name: mdm19, ID: 0x26ee566356362451, Manager
IPs: 10.103.110.19, 192.168.1.19, Management IPs:
10.103.110.19, Port: 9011
Name: mdm57, ID: 0x13c925450656db74, Tie Breaker
IPs: 10.103.110.57, 192.168.1.57, Port: 9011
```

The server *mdm57* is now a standby MDM, so it can be promoted to the MDM cluster.

10. Switch to 5-node cluster by adding the standby MDMs to the cluster:

```
scli --switch_cluster_mode --cluster_mode 5_node
--add_slave_mdm_name mdm19 --add_tb_name mdm57
```

The following output is displayed:

```
Successfully switched the cluster mode.
```

11. To view the result of the command, run:

```
scli --query_cluster
```

Output similar to the following is displayed:

```
Cluster:
Mode: 5_node, State: Normal, Active: 5/5, Replicas: 3/3
Master MDM:
Name: mdm17, ID: 0x5d07497754427fd0
```

```
IPs: 10.103.110.17, 192.168.1.17, Management IPs:
10.103.110.17, Port: 9011
Version: 2.0.972
Slave MDMs:
Name: mdm18, ID: 0x5843c4d16d8f1082
IPs: 10.103.110.18, 192.168.1.18, Management IPs:
10.103.110.18, Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm19, ID: 0x26ee566356362451
IPs: 10.103.110.19, 192.168.1.19, Management IPs:
10.103.110.19, Port: 9011
Status: Normal, Version: 2.0.972
Tie-Breakers:
Name: mdm20, ID: 0x6dfe1c5f4062b5b3
IPs: 192.168.1.20, 10.103.110.20, Port: 9011
Status: Normal, Version: 2.0.972
Name: mdm57, ID: 0x13c925450656db74
IPs: 10.103.110.57, 192.168.1.57, Port: 9011
Status: Normal, Version: 2.0.972
```

12. When changing an MDM IP address, it is mandatory to update and restart the all the SDCs in the system as well.

    a. Update the IP addresses:

    Windows:

    ```
    C:\Program Files\emc\scaleio\sdc\bin\drv_cfg --mod_mdm_ip
    --ip <EXISTING_MDM_IP_ADDRESS>
    --new_mdm_ip <NEW_MDM_IP_ADDRESSES>
    ```

    Linux:

    ```
    /opt/emc/scaleio/sdc/bin/drv_cfg --mod_mdm_ip
    --ip <EXISTING_MDM_IP_ADDRESS>
    --new_mdm_ip <NEW_MDM_IP_ADDRESSES>
    ```

    b. Restart the SDC.

    c. Verify the changes:

    Windows:

    ```
    C:\Program Files\emc\scaleio\sdc\bin\drv_cfg --query_mdms
    ```

    Linux:

    ```
    /opt/emc/scaleio/sdc/bin/drv_cfg --query_mdms
    ```

    Output similar to the following should appear:

    ```
    Retrieved 1 mdm(s)
    MDM-ID 043925027bbed30e SDC ID 28c5479b00000000
    INSTALLATION ID
    7214f7ca647c185b IPs [0]-9.4.4.12 [1]-9.4.4.11
    ```

Configure virtual IP addresses

Configure virtual IP addresses for the MDMs in your ScaleIO system.

You can configure virtual IP addresses during deployment or post-deployment. Use the following management tools to configure virtual IP addresses:

Management tool	Actions	Notes
Installation Manager (IM)	Add virtual IP addresses only.	For details, see the deployment documentation.
vSphere Web plug-in	Add virtual IP addresses only.	For details, see the *ScaleIO User Guide*.
CLI	Add, modify, and remove virtual IP addresses.	For details, see the *ScaleIO CLI Reference Guide*.
REST API	Add, modify, and remove virtual IP addresses.	For details, see REST API Reference on page 379.

Managing SDC access to the MDM

To harden SDC access to the MDM, it is possible to restrict access, pending approval of the SDC by the system. The default system setting is full access (restricted SDC mode is disabled). When the restricted SDC mode is enabled, volumes can only be mapped to "approved" SDCs. Approval is obtained by issuing the `--add_sdc` command for each SDC. You can set restricted mode before or after SDCs have been added to your network.

You can use the following commands:

Action	Command
Enable or disable restricted SDC mode	`set_restricted_sdc_mode command`
Add an SDC to the approved list, when restricted SDC mode is enabled	`--add_sdc`

For more information, see the *ScaleIO CLI Reference Guide*.

Add another IP address subnet to an MDM cluster

Add an IP network to an existing MDM cluster.

Before you begin

This topic explains how to add another IP address subnet for use by the MDM cluster. This procedure addresses scenarios where the MDM cluster uses a single network, or when an existing network needs to be replaced by a different one.

Note

This procedure describes an example for a for 3-node cluster, however, the procedure for a 5-node cluster is similar.

Procedure

1. Query the system to get the current cluster state/health:

```
scli --query_cluster
```

   Cluster status is returned, where you can identify the Master, the Slave, and the Tie Breaker.

2. Switch to single cluster mode:

```
scli --switch_cluster_mode --cluster_mode 1_node --
remove_slave_mdm_id <mdm_slave_id> --remove_tb_id <tb_id>
```

3. Remove the standby MDM:

```
scli --remove_standby_mdm --remove_mdm_id <mdm_slave_id>
```

4. Remove the Tie Breaker:

```
scli --remove_standby_mdm --remove_mdm_id <tb_id>
```

5. Add the MDM as standby with its IP addresses (including the additional IP addresses):

```
scli --add_standby_mdm --new_mdm_ip ip_1<,ip_2,...> --
mdm_role manager --new_mdm_management_ip ip_1<,ip_2,...> --
allow_asymmetric_ips --force_clean
```

   For example:

```
scli --add_standby_mdm --new_mdm_ip 10.89.9.6,10.89.11.6 --
mdm_role manager --new_mdm_management_ip 10.89.9.6,10.89.11.6
--allow_asymmetric_ips --force_clean
```

6. Add the Tie Breaker as standby with its IP addresses (including the additional IP addresses):

```
scli --add_standby_mdm --new_mdm_ip ip_1<,ip_2,...> --
mdm_role tb --new_mdm_management_ip ip_1<,ip_2,...> --
allow_asymmetric_ips --force_clean
```

7. Switch cluster operation back to a 3-node cluster:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_id <slave_id> --add_tb_id <tb_id>
```

For example:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_id 0x4520631c7262bbf1 --add_tb_id
0x3cde0ef516f61162
```

8. Query the system to get the current cluster state/health.

```
scli --query_cluster
```

Cluster status is returned, where you can check that the cluster is configured and operating as expected.

9. Switch MDM ownership to verify cluster functionality:

```
scli --switch_mdm_ownership --new_master_mdm_id
<new_master_mdm_id>
```

For example:

```
scli --switch_mdm_ownership --new_master_mdm_id
0x4520631c7262bbf1
```

10. Query the system to get the current cluster state/health.

```
scli --query_cluster
```

Cluster status is returned, where you can check that the cluster is operating as expected.

11. Add IP addresses for the Master MDM (presently Slave MDM) by following steps 2, 3, 5, 7, and 8.
12. Optional: Switch MDM ownership back to the original MDM:

```
scli --switch_mdm_ownership --new_master_mdm_id MDM_ID
```

Managing the SDSs and cache

You can modify and SDS port while there is I/O running. You can also manage the system's cache, including Read Flash cache and Read RAM cache.

Modifying an SDS port during IO

If you need to modify an SDS port (on-the-fly) while there is IO running, perform the following steps:

Note

Sometimes, your network topology needs to be prepared for the addition of new port, and this may take some time. The ScaleIO system does not prevent unnecessary degraded status or disconnection of SDS. Therefore, in such situations, it is recommended to place the SDS in Maintenance Mode before commencing this procedure.

Procedure

1. On the SDS, perform one of the following, depending on the Operating System:

Operating System	Task
Linux	Run the script: `/opt/emc/scaleio/sds/bin/close_firewall_port.sh`
Windows	From command line, run the batch file `C:\Program Files\EMC\scaleio\sds\bin\close_firewall_port.bat`

2. Open the following SDS file with a text editor, and change the port number shown there to the new port number:

Operating System	File name
Linux	`/opt/emc/scaleio/sds/bin/port`
Windows	`C:\Program Files\EMC\scaleio\sds\bin\port`

3. Open the following SDS configuration file with a text editor:

Operating System	File name
Linux	`/opt/emc/scaleio/sds/cfg/conf.txt`
Windows	`C:\Program Files\EMC\scaleio\sds\cfg\conf.txt`

4. Add the parameter `tgt_port = <NEW_PORT_NUM>` to the file, where <NEW_PORT_NUM> represents the new port number.
5. Perform one of the following:

Operating System	Task
Linux	Run the script: `/opt/emc/scaleio/sds/bin/open_firewall_port.sh`
Windows	From command line, run the batch file `C:\Program Files\EMC\scaleio\sds\bin\open_firewall_port.bat`

6. On the SDS, perform one of the following:

Operating System	Task
Linux	Run the command: `Pkill sds`
Windows	From command line, run the command: `net stop sds_service & net start sds_service`

7. On the MDM, modify the SDS port using the command:

```
scli --modify_sds_port (--sds_id <ID> | --sds_name <NAME> | --
sds_ip <IP>) --new_sds_port <PORT>
```

For example, for an SDS called "sds198" where the new port number is 7071, type:

```
scli --modify_sds_port --sds_name sds198 --new_sds_port 7071
```

Note

If you modify the SDS port on the MDM first, instead of following the above procedure, IO errors might be encountered.

Managing Read Flash cache

This section describes how to manage the Read Flash Cache (RFcache) feature, which uses PCI flash cards and SSDs for caching of the HDDs in the SDS, thus accelerating the reads of its HDD devices.

RFcache devices and configuration can be performed during initial system deployment. If you want to add the use of the RFcache feature after deployment, use the following work flow:

1. Ensure that the RFcache policy is enabled in the Storage Pool.
2. Enable RFcache in the SDSs where RFcache devices will be added.
3. Add the RFcache devices to the SDSs.

   If you want to stop using the RFcache feature, or remove a specific RFcache device, use the following work flow:
4. Stop RFcache usage in the Storage Pool
5. Remove (command) the RFcache device from the SDS.
6. You can then do one of the following:
   - Physically remove the device from the chassis, and then restart RFcache usage in the Storage Pool
   - Add the device to the SDS as a storage device

For more information, see Command Quick Reference on page 99.

Managing RAM read cache

This section describes how to manage the RAM read cache feature, which is designed to allocate RAM on nodes for caching of reads or writes. The feature is configured at the following levels:

- For a volume (optional)—specific volume to use or not use caching. If you want all I/Os for a specific volume to use caching, make sure that the volume, the corresponding Storage Pool, and its SDSs are all configured for using caching.
- For one or more Storage Pools—caching must be configured at this level for caching to work in the corresponding SDSs. When Storage Pools are created, caching is not automatically enabled by default, unless the `--use_rmcache` option is added to the `add_storage_pool` command. Write handling mode can also be configured via the `add_storage_pool` command, or in a separate command. To run CLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *ScaleIO CLI Reference Guide*.
- For individual SDSs—caching may be disabled at this level, even if caching in the corresponding Storage Pool is enabled. Cache size can be configured at the SDS level.

  RAM Read Cache is disabled by default on Volumes and Storage Pools. SDSs, on the other hand, are enabled by default, but the Storage Pool setting overrides the SDS setting.

  The amount of RAM you can allocate for cache is limited by the amount of RAM on the SDS server:

- If the RAM is less than 32 GB, 50% of memory can be used for cache
- If the RAM is more than 32 GB, 75% of memory can be used for cache

  The maximum amount of RAM cache is described in Table 3 on page 32.

Note

Only blocks up to 128 k in size will be cached. Any blocks larger than 128k will be ignored by caching.

For a read to be stored in a specific SDS cache, you have to make sure that the cache on that SDS is enabled, and the relevant Storage Pool and the relevant volume are both configured to use cache.

For more information, see Command Quick Reference on page 99.

SDC operations

Many SDC operations use drv_cfg. The drv_cfg command line is a local CLI utility that affects only the client on which the SDC is running. Possible SDC operations include updating the SDC driver with IP changes, detecting new volumes, querying volumes, loading a configuration file, adding an MDM, modifying an MDM IP address, enabling support of PDL state, and more.

Updating the SDC driver with IP changes

Procedure

1. Edit `drv_cfg.txt` and change the IP address in the last line to the new IP.

Location of `drv_cfg.txt`:

- Linux: `/bin/emc/scaleio/`
- Windows: In the following registry key - `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\scini\Parameters\mdms`

Enter the IP addresses as a comma-separated list.

Note

On ESXi, GUID and MDM lists are stored as module parameters, and not in a `drv_cfg.txt` file. To modify these parameters, use `esxcli` commands.

---

2. Save and close the file.
3. Type the following command:

```
/etc/init.d/scini restart
```

Detecting new volumes

Command

`drv_cfg --rescan`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --rescan
```

Description/Notes

Volumes are always exposed to the operating system as devices with the prefix `scini` (such as `/dev/scinia`, `/dev/scinib` and so on). Unique names can be found under `/dev/disk/by-id/`.

ScaleIO periodically scans the system to detect new volumes. You can initiate a scan for the most up-to-date status on a particular SDC node. This command is unique because it is not a CLI command, but rather a command issued on the specific SDC.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESX: Contact Customer Support for access to this tool.

For further details on how to set the mounting options see Mounting ScaleIO on page 286.

Parameters

Not applicable.

Query volumes using drv_cfg

Command

`drv_cfg --query_vols`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_vols
```

Description/Notes

This utility retrieves information about all known active volume objects in kernel mode. You can use this utility to determine which volumes are mapped, and the ID of each volume in the ScaleIO system.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESX: Contact Customer Support for access to this tool.

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_vols
```

Query tgt objects using drv_cfg

Command

`drv_cfg --query_tgts`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_tgts
```

Description/Notes

This utility retrieves information about all known active tgt objects in kernel mode.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESX: Contact Customer Support for access to this tool.

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_tgts
```

Query GUID using drv_cfg

Command

`drv_cfg --query_guid`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_guid
```

Description/Notes

This utility retrieves the unique ID of the kernel module. The utility can be used to verify that all SDC GUIDs in the system are unique.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESX: Contact Customer Support for access to this tool.

Note

If the SDC was removed and reinstalled, the GUID of the SDC will be different to its original GUID. In such a case, you may need to remove the SDC, if two SDCs now have the same GUID.

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_guid
```

Query MDMs using drv_cfg

Command

`drv_cfg --query_mdms`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_mdms
```

Description/Notes

This utility retrieves information about all known MDM objects in kernel mode. This utility is typically used to determine to which MDM an SDC is connected.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESX: Contact Customer Support for access to this tool.

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --query_mdms
```

Loading a configuration file using drv_cfg

Command

`drv_cfg --load_cfg_file`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

This command can not be used on ESX servers. Instead, follow the steps described in "Modifying configuration parameters on ESX servers".

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg
--load_cfg_file <FILE_NAME>
```

Description/Notes

This utility reads a configuration file containing MDM IP addresses, and calls the kernel to connect to them.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESX: Contact Customer Support for access to this tool.

The configuration file that is loaded when using the `drv_cfg --load_cfg_file` utility is not persistent; when you restart the SDC, the changes will be lost.

To make the changes persistent, perform either of the following:

- Install the SDC on every server that will expose ScaleIO volumes to the application running, by executing the following command:

```
MDM_IP=<IP of the MDM> rpm -i <full rpm file path>
```

- Use the following `drv_cfg` command:

```
/opt/emc/scaleio/sdc/bin/drv_cfg --mod_mdm_ip
--ip <EXISTING_MDM_IP_ADDRESS> --new_mdm_ip
<NEW_MDM_IP_ADDRESSES>
```

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg
--load_cfg_file /bin/emc/scaleio/drv_cfg.txt
```

Adding an MDM using drv_cfg

Command

`drv_cfg --add_mdm`

Note

This is not a CLI command, but rather an executable that is run on the SDC server.

This command can not be used on ESX servers. Instead, follow the steps described in "Modifying configuration parameters on ESX servers".

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --add_mdm
--ip <MDM_IP_ADDRESS_LIST>
```

Description/Notes

This utility calls the kernel module to connect to an MDM. This command is typically used in cases where an SDC is connected to more than one ScaleIO system.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESX: Contact Customer Support for access to this tool.

Note

Extending your ScaleIO system with another MDM requires that you update all SDCs in your system with the new MDM IP address. Run the `drv_cfg` utility with the `--mod_mdm_ip` option (see "Modifying an MDM IP address using drv_cfg"), and to make the change persistent, use the `--file` parameter. In addition, any additional objects or systems which interface with the MDM must also be updated. For more information, see "Modifying an MDM's management IP address" in the *ScaleIO CLI Reference Guide*.

Parameters

Parameter	Description
--ip <MDM_IP_ADDRESS_LIST>	List of IP addresses (comma delimited) for this Master or Slave MDM
Optional:	
--file <CONFIG_FILE_NAME>	Name of the configuration file to which the MDM information should be written
--only_cfg	Do not call the kernel to actually connect

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --add_mdm
--ip 10.100.22.20,10.100.22.30
--file /bin/emc/scaleio/drv_cfg.txt
```

Modifying an MDM IP address using drv_cfg

Command

`drv_cfg --mod_mdm_ip`

Note

This is not a CLI command, but rather an executable that is run on the SDC server. This command can not be used on ESX servers. Instead, follow the steps described in "Modifying configuration parameters on ESX servers".

Syntax

```
/opt/emc/scaleio/sdc/bin/drv_cfg --mod_mdm_ip
--ip <EXISTING_MDM_IP_ADDRESS>
--new_mdm_ip <NEW_MDM_IP_ADDRESSES> [--file <CONFIG_FILE_NAME>]
["-""-"only_cfg]
```

Description/Notes

This utility calls the kernel to modify an MDM's IP address list. It is typically used in cases when an MDM IP address has changed, or when MDMs are added/removed

from/to the system. The command must be run on every SDC in the system. To bring the changes into effect, a server restart is required.

Location of `drv_cfg` command:

- Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
- Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
- ESX: Contact Customer Support for access to this tool.

Note

Extending your ScaleIO system with another MDM requires that you update all SDCs in your system with the new MDM IP address. Ensure that you run this command on all SDCs. For more information, see the last step in Replacing a cluster member without adding a new server to the cluster on page 124 .

Parameters

Parameter	Description
--ip <EXISTING_MDM_IP_ADDRESS>	One of the existing MDM IP addresses
--new_mdm_ip <NEW_MDM_IP_ADDRESSES>	The new IP address list (comma delimited) for this MDM. If you want to retain the existing address(es), include them in this list.
Optional:	
--file <CONFIG_FILE_NAME>	The name of the configuration file to which the MDM information should be written
--only_cfg	Do not call the kernel to actually connect

Example

```
/opt/emc/scaleio/sdc/bin/drv_cfg --mod_mdm_ip
--ip 10.100.20.20
--new_mdm_ip 10.100.20.20,10.100.20.30,10.100.20.40
```

Permanent Device Loss state

When the MDM has disconnected from the SDC and a volume mapped to this SDC has experienced an I/O error, the ESXi host continuously sends I/Os to the device to determine if the device has become accessible. This can subsequently cause a high I/O-error load that can lead to the host freezing. In cases where the device is disconnected long-term, such as when the entire MDM cluster is down during a network upgrade, the SDC can change the volume state to Permanent Device Loss (PDL) to prevent more I/O errors coming from the ESXi.

PDL is an ESXi state, which, once enabled, is supported in ScaleIO. Once the ESXi host loses connectivity with a device, if a timeout value is reached, the ESXi will be notified that the device is in a PDL state. The timeout value can be manually set. Once a device is in a PDL state, the ESXi host no longer attempts to re-establish connectivity or issue commands to the device.

Recovering a device from PDL state is described in the VMware documentation for your operating system version. The following link is for ESXi v6.5: https://

docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.storage.doc/GUID-A513D44C-71DE-47ED-B781-327F78659404.html

Use the instructions that match your environment.

Enabling support of PDL state on the ESXi

Enable support of Permanent Device Loss (PDL) state on the ESXi host.

Procedure

1. On the ESXi host:

```
esxcli system module parameters set -m scini
"<<PREVIOUS_MODULE_PARAMS>> bBlkDevIsPdlActive=1
blkDevPdlTimeoutMillis=<TIMEOUT_VALUE>"
```

   where:

   - *<<PREVIOUS_MODULE_PARAMS>>* is any previous module parameters being used for this ESXi host.
   - *<TIMEOUT_VALUE>* is the timeout time in milliseconds. Its value can be 1000-3600000 (default is 60000) and including it in the command is optional.

2. Reboot the host.

Disabling support of PDL state on the ESXi

Disable support of Permanent Device Loss (PDL) state on the ESXi host.

Procedure

1. On the ESXi host:

```
esxcli system module parameters set -m scini
"<<PREVIOUS_MODULE_PARAMS>> bBlkDevIsPdlActive=0"
```

   where *<<PREVIOUS_MODULE_PARAMS>>* is any previous module parameters being used for this ESXi host.

2. Reboot the host.

Managing ESX servers

Using the following procedures, you can modify parameters on ESX servers and check the SDC state on ESX servers.

Modifying parameters on ESX servers

On an SDC running on an ESX server, esxcli commands can be used in the following cases:

- MDM IP addresses need to be added to the existing list on an SDC
- MDM IP addresses need to be replaced on an SDC
- The SDC's GUID needs to be changed

Specifically, the SDC's GUID or IP address needs to be identified, and then used to add or modify the MDM IP addresses or GUID (depending on the parameter that you

want to modify). If you want to add additional MDM IP addresses to existing ones, you must list both old and additional IP addresses in the esxcli command.

For more information about SDC tuning, see *ScaleIO Performance Fine-Tuning Technical Notes*.

If the current configuration of the ScaleIO system is registered with a v2.0 VMware plug-in, you can use the plug-in Update SDC parameters to update the MDM IP addresses. For more information, see the *EMC ScaleIO Deployment Guide*.

Note

These procedures require a server restart to apply the new configuration. The configuration will remain persistent after future server restarts.

To configure MDM IP addresses on the SDC, perform these steps:

Procedure

1. Find the SDC's GUID and the MDM IP addresses configured on the ESX, by typing the command:

```
esxcli system module parameters list -m scini
```

2. In the output of the command, find the existing GUID and MDM IP addresses.

   For example, in the output excerpt below, the GUID and IP addresses are marked in bold:

```
IoctlIniGuidStr string
39b89295-5cfc-4a42-bf89-4cc7e55a1e5b Ini Guid, for example:
12345678-90AB-CDEF-1234-567890ABCDEF
```

```
IoctlMdmIPStr string
9.99.101.22,9.99.101.23 Mdms IPs, IPs for MDM in same cluster
should be comma-separated. To configure more than one cluster
use '+' to separate between IPs.For Example:
10.20.30.40,50.60.70.80+11.22.33.44. Max 1024 characters
```

3. To configure the MDM IP addresses on the SDC, type the command

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=<GUID> IoctlMdmIPStr=<MDM_IPS>"
```

   where *<GUID>* is the existing SDC GUID that you identified in the previous step, and *<MDM_IPS>* is the list of MDM IP addresses. A maximum of 1024 characters is allowed.

   a. To replace the old MDM IP addresses with new MDM IP addresses, omit the old addresses from the command.

   b. To add MDM IP addresses to the existing IP addresses, type both the existing IP addresses and the new IP addresses in the command.

MDM IP addresses for MDMs in same cluster must be comma-separated. To configure more than one cluster, use '+' to separate between IP addresses in different clusters. For example:

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=39b89295-5cfc-4a42-bf89-4cc7e55a1e5b
IoctlMdmIPStr=10.20.30.40,50.60.70.80+11.22.33.44"
```

4. To apply the new configuration, restart the ESX server.

   To change the GUID of the SDC, perform these steps:

5. Find the SDC's GUID and the MDM IP addresses configured on the ESX, by typing the command

```
esxcli system module parameters list -m scini
```

6. In the output of the command, find the existing GUID and MDM IP addresses.

   For example, in the output excerpt below, the GUID and IP addresses are marked in bold:

```
IoctlIniGuidStr string
39b89295-5cfc-4a42-bf89-4cc7e55a1e5b Ini Guid, for example:
12345678-90AB-CDEF-1234-567890ABCDEF
```

```
IoctlMdmIPStr string 9.99.101.22,9.99.101.23 Mdms IPs, IPs for
MDM in same cluster should be comma-separated. To configure
more than one cluster use '+' to separate between IPs.For
Example:
10.20.30.40,50.60.70.80+11.22.33.44. Max 1024 characters
```

7. To change the GUID on the SDC, type the command

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=<NEW_GUID> IoctlMdmIPStr=<MDM_IPS>
```

   where <NEW_GUID> is the new SDC GUID, and <MDM_IPS> is the list of MDM IP addresses that you identified in the previous step. You must include these IP addresses in the command.

   For example:

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=28a78184-4beb-4a42-bf89-4cc7e55a1e5b
IoctlMdmIPStr= 9.99.101.22,9.99.101.23"
```

8. To apply the new configuration, restart the ESX server.

Checking the SDC state on ESX servers

On an SDC running on an ESX server, an esxcli command can be used to check the current state of the SDC.

To display the SDC state on the ESX server, type the following command:

```
esxcli system module list |grep scini
```

The following examples show typical outputs of the command:

- Output where driver is installed and enabled to load, but not loaded:

```
Name Is Loaded Is Enabled
---------------------------- --------- ----------
scini false true
```

- Example of SDC in correct state (enabled and loaded):

```
Name Is Loaded Is Enabled

---------------------------- --------- ----------
scini true true
```

CHAPTER 8

Security Management

This chapter describes security management in ScaleIO. Topics include:

- Setting up SSH authentication on the ScaleIO Gateway 148
- Configuring SSL component authentication 148
- Managing SDC access to the MDM 152
- Approved encryption methods 153
- Login banner overview 153

Setting up SSH authentication on the ScaleIO Gateway

A manually generated public-private key pair can be used to perform SSH key authentication, instead of passwords, between the ScaleIO Gateway and ScaleIO system servers. For more information, see "Using SSH authentication on the ScaleIO Gateway" in the ScaleIO Deployment Guide.

Configuring SSL component authentication

ScaleIO uses SSL authentication to authenticate both internal system components, and communication between the MDM and external components such as the ScaleIO Gateway, GUI clients, vSphere plug-in, and CLI clients. Secure communication is typically installed and configured by default during system deployment.

Note

If your system has been upgraded from a version earlier than version 2.0, or if secure communication between components was disabled during installation, follow the instructions provided in the section "Switching to secured authentication mode" in the ScaleIO Deployment Guide.

Internal component authentication

When this feature is enabled, the MDM generates a self-signed certificate for itself, and the SDSs generate certificates signed by the MDM's certificate. The MDM has a single certificate for the entire cluster. The certificate is stored in the MDM repository.

Each SDS has its own SSL certificate file:

- Linux: `/opt/emc/scaleio/sds/cfg/sds_certificate.pem`
- Windows: `C:\Program Files\emc\scaleio\sds\cfg\sds_certificate.pem`

  When an SDS is added to the cluster, the MDM receives a CSR (Certificate Signing Request) from the SDS, signs it with its own internal certificate and returns it to the SDS to be stored in its local key-store. If the SDS disconnects and reconnects, the MDM must authenticate it.

External component authentication

Secure communications can be performed between the MDM and the following external components, and are typically enabled during deployment of the system:

- ScaleIO Gateway—The ScaleIO Gateway maintains the SSL certificates for itself and for the following components:
  - SNMP
  - REST API
  - IM
- vSphere plug-in
- GUI

- CLI

Workflow for self-signed security certificates

The system generates and signs self-signed certificates automatically when secure communication is enabled, and no user intervention is required. If you want to replace these certificates with new self-signed ones, follow this workflow:

Procedure

1. Run the command `scli --generate_mdm_certificate`.

   To run CLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *ScaleIO CLI Reference Guide*.

2. When using the CLI, on the first connection to the MDM, the CLI will display the MDM's certificate and will prompt the user to approve the certificate.

   Upon approval, the trusted certificate will be saved.

3. When using the GUI, approve the MDM certificate at login, and then approve other certificates using the **Preferences** menu, **Renew Certificates** option.

Workflow for externally signed security certificates

The system generates and signs self-signed certificates automatically when secure communication is enabled, and no user intervention is required. If you want to replace these certificates with ones signed by an external Certificate Authority, follow this workflow:

Procedure

1. Log in to the system using the `scli --login` command as either a root user (on Linux) or as an administrator (on Windows).
2. Generate a CSR file, using the command `scli --generate_mdm_csr_file --target_mdm_ip <IP_ADDRESS>`.

   A file called `mdm-target_hostname.csr` will be created in the location:

   a. Linux: `/opt/emc/scaleio/mdm/cfg`

   b. Windows: `C:\Program Files\emc\scaleio\mdm\cfg`

   To run CLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *ScaleIO CLI Reference Guide*.

3. Submit the CSR file created in the previous step to your Certificate Authority.

   The Certificate Authority must sign your CSR and return two files to you:

   a. Certificate for your MDM

   b. Certificate Authority “Trusted” or “Root” certificate

4. Save the signed certificate for the MDM in the location:

   a. Linux: `/opt/emc/scaleio/mdm/cfg`

   b. Windows: `C:\Program Files\emc\scaleio\mdm\cfg`

5. Manually change the MDM certificate’s file name to `mdm_signed_certificate.pem`.

6. Run the following script on the directory:

```
./apply_signed_certificate.py --mdm_ip <IP_address> --
local_mdm_ip <IP_address>
```

where `--mdm_ip` is the IP address of the Master MDM, and `--local_mdm_ip` is the IP address of the MDM where you want to change the certificate.

If the remote read-only feature is enabled on the MDM, add `--skip_cli_command` to the command, and later, while logged in with security permissions, run the command `scli --replace_mdm_security_files`.

Note

This step changes the MDM certificate, and might cause a brief single point of failure period (switch ownership).

7. For all external components that will communicate with the MDM (GUI, CLI, vSphere Plugin, REST, IM) add the Trusted or Root certificate from the Certificate Authority to each component.

   The Trusted/Root certificate must be added to the file called `truststore.jks`, using Keytool.

   For more information, see Using Keytool to add certificates to external components on page 150.

8. When using the CLI, on the first connection to the MDM, the CLI will display a message similar to the following:

```
[root@112CC-4~]# scli --login --username admin --password
Scaleio018 Certificate required for issuer: /C=US/ST=MA/
L=Hopkinton/O=EMC-Scaleio1213/CN=Scaleio018
Please add the certificate with scli --add_certificate
```

   Add the Trusted/Root certificate using the `--add_certificate` command. For more information, see the *ScaleIO CLI Reference Guide*.

Using Keytool to add certificates to external components

This topic explains how to add Certificate Authority certificates to ScaleIO external components. The `truststore.jks` file located on all components saves all the MDM/LIA certificates approved by the client. The file's location depends on the management client and operating system:

Gateway

- Linux:

  `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/certificates`

- Windows (64-bit):

  `C:\Program Files\EMC\ScaleIO\Gateway\webapps\ROOT\WEB-INF\classes\certificates`

GUI

- Windows:

  `C:\Users\[user_name]\AppData\Roaming\EMC\scaleio\certificates`

vSphere

- Linux:

  `$HOME/.vmware/scaleio/certificates`

- Windows:

  `C:\Users\[user_name]\AppData\Roaming\VMware\scaleio\certificates\truststore.jks`

  `C:\Windows\System32\config\systemprofile\AppData\Roaming\VMware\scaleio\certificates`

Using Keytool

Use the Java Keytool utility to modify or view the content of the trust store file. The remainder of this topic lists some useful Keytool commands. Keytool is a part of the Java (JRE or JDK) installation and can be found in the `bin` directory.You can add -storepass changeit to all commands that require a password. The password for the trust store is "changeit" (Java default).

Note

The certificate alias must be unique in the trust store file. We usually use the certificate's full subject.

For example: givenname=mdm, ou=asd, o=emc, l=hopkinton, st=massachusetts, c=us, cn=centos-6.4-adi5

- List the certificates in the trust store:

  ```
  keytool -list -v -keystore [path_to_certificates_folder]/
  truststore.jks
  ```

  Example:

  ```
  keytool -list -v -keystore C:\Users\cj\AppData\Roaming\EMC
  \scaleio\certificates\truststore.jks
  ```

- Check a particular entry using an alias:

  ```
  keytool -list -v -keystore [path_to_certificates_folder]/
  truststore.jks -alias [unique_alias] -storepass changeit
  ```

  Example:

  ```
  keytool -v -list -keystore C:\Users\cj\AppData\Roaming\EMC
  \scaleio\certificates\truststore.jks \truststore.jks -alias
  ```

```
"givenname=mdm, ou=asd, o=emc, l=hopkinton, st=massachusetts,
c=us, cn=centos-6.4-adi5"
```

- Add a new trusted certificate to the trust store:

```
keytool -import -trustcacerts -alias [unique_alias] -file
[path_to_the_certificate_file] -keystore
[path_to_certificates_folder]/truststore.jks
```

Example:

```
keytool -import -trustcacerts -alias "givenname=mdm, ou=asd,
o=emc, l=hopkinton, st=massachusetts, c=us, cn=centos-6.4-adi5"
-file c:\temp\centos-6.4-adi5.cer -keystore C:\Users\cj\AppData
\Roaming\EMC\scaleio\certificates\truststore.jks
```

- Delete a certificate from the trust store:

```
keytool -delete -alias [unique_alias] -keystore
[path_to_certificates_folder]/truststore.jks
```

Example:

```
keytool -delete -alias "givenname=mdm, ou=asd, o=emc,
l=hopkinton, st=massachusetts, c=us, cn=centos-6.4-adi5" -
keystore C:\Users\cj\AppData\Roaming\EMC\scaleio\certificates
\truststore.jks
```

- Export a certificate from the trust store:

  `keytool -export -alias [unique_alias] -file [certificate_file_path] -keystore [path_to_certificates_folder]/truststore.jks`

  Example:

```
keytool -export -alias "givenname=mdm, ou=asd, o=emc,
l=hopkinton, st=massachusetts, c=us, cn=centos-6.4-adi5" -file
c:\temp\centos-6.4-adi5.cer -keystore C:\Users\cj\AppData
\Roaming\EMC\scaleio\certificates\truststore.jks
```

Managing SDC access to the MDM

To harden SDC access to the MDM, it is possible to restrict access, pending approval of the SDC by the system. The default system setting is full access (restricted SDC mode is disabled). When the restricted SDC mode is enabled, volumes can only be mapped to “approved” SDCs. Approval is obtained by issuing the `--add_sdc` command for each SDC. You can set restricted mode before or after SDCs have been added to your network.

You can use the following commands:

Action	Command
Enable or disable restricted SDC mode	`set_restricted_sdc_mode command`
Add an SDC to the approved list, when restricted SDC mode is enabled	`--add_sdc`

For more information, see the *ScaleIO CLI Reference Guide*.

Approved encryption methods

A specific set of encryption methods are approved for use with your system.

The following encryption methods are approved for use:

- TLS_DHE_DSS_WITH_AES_128_CBC_SHA256
- TLS_DHE_DSS_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256
- TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
- TLS_RSA_WITH_AES_256_CBC_SHA256

Note

In order to use CURL on RHEL6 with ScaleIO Gateway v2.0.0.3 and higher, upgrade the NSS package to 3.21.0. (use the YUM update command).

Login banner overview

A login banner is a text file that is displayed upon login to the system. It can be used to communicate messages or to obtain user consent to real-time monitoring of information and retrieval of stored files.

When the login banner is set up, it appears during the system login process before the login credential prompts. The login banner displays differently in the ScaleIO GUI and CLI interfaces:

- GUI:
  - When logging in, the login banner is displayed, and must be approved.
- CLI:
  - When logging in, the user is prompted to press any key, after which the banner is displayed.
  - To continue, the banner must be approved.

Limitations:

- Only users with administrative security rights can set up, update, or remove the login banner.
- Supported in Windows and RHEL operating systems.
- Text files up to 16 bytes are supported.
- Only one login banner is supported.

Setting up a login banner using the CLI

You can use the CLI to set up, modify, or stop displaying a login banner.

Before you begin

Ensure that you have access to the IP address of the Master MDM.

Procedure

1. Log in to the ScaleIO system using the IP address of the Master MDM.
2. Perform the desired operation:

Option	Description
Create (or modify) a new banner	a. Create a text file (or modify an existing file) with the message that you want to display in the login banner. b. Run the following command: `scli --set_login_banner --filename <FILENAME>` where *<FILENAME>* is the path of the login banner text file. The login banner is displayed the next time a user logs into the ScaleIO system.
Stop displaying the banner	a. Run the following command: `scli --set_login_banner --remove_banner`

Enabling or disabling preemptive acceptance of the login banner

Preemptive acceptance of the login banner allows the user to bypass the login banner, for example, when running scripts. A user with admin security rights can enable or disable the option of preemptive acceptance. By default, preemptive acceptance is enabled and the login banner can be bypassed using a CLI command.

Before you begin

To enable or disable the preemptive acceptance option, you must have administrative rights.

Procedure

1. Log in to ScaleIO:

```
scli --login --username admin --password <PASSWORD>
```

2. Run the following command to enable preemptive acceptance:

```
scli --set_cli_login_banner_preemptive_acceptance --enable
```

3. Run the following command to disable preemptive acceptance:

```
scli --set_cli_login_banner_preemptive_acceptance --disable
```

Activating preemptive acceptance of the login banner

When preemptive acceptance of the login banner is enabled (default), you can log in to ScaleIO in a special way that activates preemptive acceptance of the login banner.

Before you begin

Preemptive acceptance of the login banner is enabled.

Procedure

1. Log in to ScaleIO with the `accept_banner_by_scripts_only` parameter:

```
scli --login --username <USERNAME> --
accept_banner_by_scripts_only
```

where *<USERNAME>* is the user running the script.

CHAPTER 9

Opening the GUI and Logging In

This chapter describes how to open the GUI, and log in procedures. Topics include:

- Log in to the ScaleIO GUI 158
- Connection and disconnection information 158

Log in to the ScaleIO GUI

Open and log in to the ScaleIO GUI.

Before you begin

Ensure that:

- The GUI software is installed on the workstation. To install the GUI, see "Installing the ScaleIO GUI."
- You have these credentials (available from the administrator):
  - MDM management IP address or hostname
  - Username (default: admin)
  - Password (defined during deployment)

Procedure

1. Open the GUI by selecting **Start** > **All Programs** > **ScaleIO GUI**.

   The initial login screen is displayed.
2. Type the IP address or hostname and click **Connect**.

   If a certificate notice is displayed, review and accept the certificate.

   If a login banner is displayed, confirm it to continue.
3. In the login screen, type the username and password, and click **Login**.

Results

The ScaleIO GUI is displayed.

After you finish

Users and passwords are configured with the ScaleIO CLI. For more information, see the "Security" chapter of the *ScaleIO User Guide*.

Connection and disconnection information

You can check at any time to which IP address your GUI is connected, using the following methods:

- View the IP address displayed in the top left corner of the GUI window.
- Hover your mouse pointer over the **Management** tile on the Dashboard. A tooltip displays connection information for the nodes in the MDM cluster, and the management IP addresses
  If your GUI loses its connection with the MDM, the window display is dimmed, and a notification dialog box is displayed.

CHAPTER 10

GUI Features

This chapter describes GUI features, and includes the following topics:

- GUI overview 160
- GUI conventions 160
- Dashboard view 162
- Frontend views 170
- Backend view 173
- Alerts view 179
- Property Sheets 180

GUI overview

This chapter describes how to use the Graphical User Interface (GUI) to monitor and configure ScaleIO.

The various windows display different views and data that are beneficial to the storage administrator. You can review the overall status of the system, drill down to the object level, and monitor these objects. You can use the GUI to provision and modify some of the objects.

The following sections in this chapter describe the available windows, and how to use them. The Glossary at the end of this publication provides more detailed information about the objects and properties displayed in the GUI. Command Quick Reference on page 99 provides information about the GUI commands that can also be performed using the CLI.

The following table provides a general overview of the tasks that you can perform with the GUI.

Table 22 GUI task overview

To do this...	Use...
See a general overview of the entire system	Dashboard view—see Dashboard view on page 162
See detailed information about one or more backend system objects, in table format (filtering available)	Backend view—see Backend view on page 173
See detailed information about one or more frontend system objects, in table format (filtering available)	Frontend views—seeFrontend views on page 170
See very detailed information about a specific system object	Property Sheets—see Property Sheets on page 180
Minimize the main window to a floating widget	Widget—see Widget on page 164
See a list of errors and alerts currently active in the system	Alerts view—see Alerts view on page 179
Add, remove, maintain, or configure a system object (backend)	Backend view—see Configuring the System using the GUI on page 187
Add, remove, or configure volumes, snapshots or SDCs	Frontend views—see Configuring the System using the GUI on page 187
Monitor various aspects of your system	Backend and Frontend views—see Monitoring the System using the GUI on page 185
Find information about your system's license	Preferences menu—see Viewing licensing information on page 186

GUI conventions

This section describes conventions used in the ScaleIO GUI, including alert indicators and color codes.

Alerts indicators

The Alerts indicators show the overall error state of the system. When lit, indicators show the number of active alerts of each severity. Similar indicators are displayed in some views of the Backend table, and also on Property Sheets (in some cases, an additional blue indicator for information only is included). You can view details about the alerts active in the system in the Alerts view. For more information about Alerts, see Alerts view on page 179 and ScaleIO Alerts in SNMP, GUI, REST, and ESRS on page 535.

Figure 15 Alerts indicators

Color codes

Color codes provide quick visual feedback on the status of various objects in the system. The following tables summarize the colors used in the system, and their meaning. The color codes are used in a variety of elements and views in the user interface.

Table 23 GUI color codes

Color	Meaning	Dashboard View	Backend View
LIGHT GREEN (protected)	Available protected, healthy storage	✓	✓
GREEN (in maintenance)	One or more SDSs are in Maintenance Mode, and part of those SDSs' capacity is temporarily protected on other SDSs	✓	✓
YELLOW	Snapshot capacity (yellow outline)	✓	
	Capacity-related statuses		✓
ORANGE (degraded)	Data is not protected	✓	✓
	Rebuild or Rebalance in progress	✓	
RED (failed)	Data is unavailable	✓	✓
DARK GRAY	Unused capacity	✓	
	No activity or zero values		✓
DARK GRAY striped with RED	System is unable to determine if capacity is Unavailable or Unused	✓	✓
PALE GRAY	Decreased capacity. This capacity exists physically, but has been disabled (typically to allow maintenance tasks on devices).	✓	

Table 23 GUI color codes (continued)

Color	Meaning	Dashboard View	Backend View
BLUE (spare)	Capacity reserved for recovery purposes	✓	
BRONZE	Volume capacity	✓	
DARK BLUE	Indicates selected items in the filter		✓

Table 24 Alert symbols and color codes

Color	Meaning	Dashboard View	Backend View	Alert View
YELLOW	Low alert status	✓	✓	✓
ORANGE	Medium alert status	✓	✓	✓
RED	High alert status	✓	✓	✓
LIGHT BLUE	Information message (no faults)		✓	✓

Dashboard view

The Dashboard displays the overall system status. Each tile displays a certain aspect of the storage system. Various controls let you customize the information displayed on the Dashboard. The following figure shows the Dashboard controls.

Figure 16 Dashboard controls

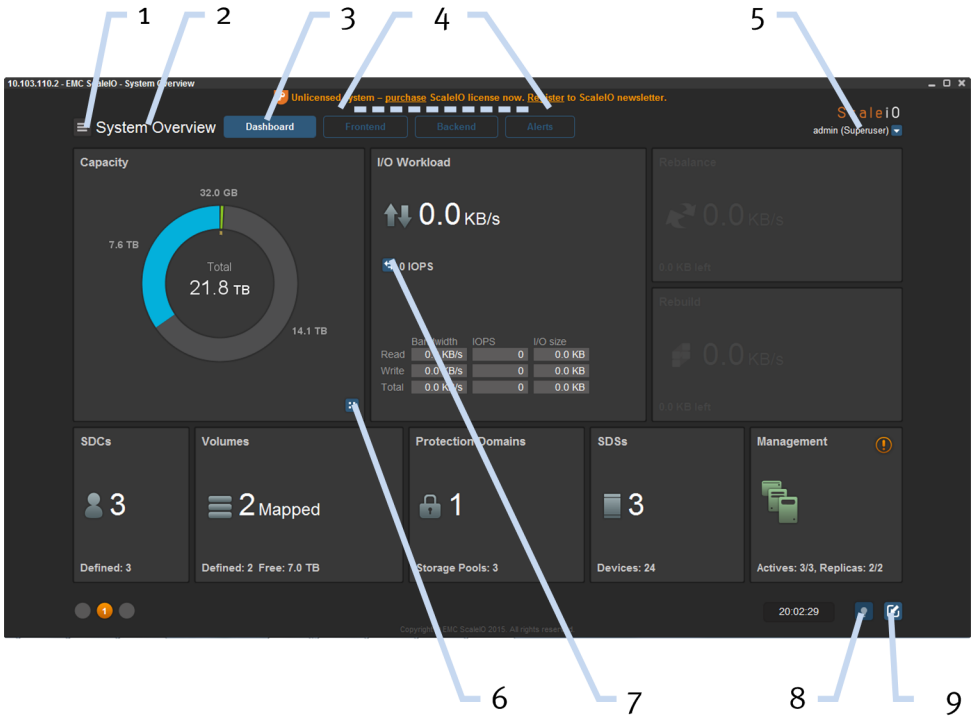

1 **Navigation toggle**—Displays navigation tree

2 **Current view**—Displays object currently being applied to the dashboard

3 **Dashboard**—Displays the Dashboard view, currently selected

4 **Frontend**, **Backend** and **Alerts** tabs —Display the corresponding views

5 **Preferences** —Settings, Add Nodes, Upgrade, Log Collection, User Preferences, Logout, About — Settings, Add Nodes, Upgrade, Log Collection, System Settings, User Preferences, Logout, About

6 **Color code legend**—Displays the legend

7 **I/O Workload toggles**—control the information displayed in this tile. For more information, see Customizing system preferences on page 223.

8 **Dimmer toggle**—Enables\disables the dimmer

9 **Widget toggle**—Minimizes the display to a widget

Navigation tree

The Dashboard's navigation button toggles the display of the navigation tree. The navigation tree is hierarchical, and controls the Dashboard display. You can display information on the Dashboard according to:

- Entire system (default)
- Protection Domain
- Storage Pool

You can change the Dashboard display by double-clicking the desired navigation tree node. Some tiles on the Dashboard may be dimmed if they are not relevant for the node that you have selected.

Preferences menu

The **Preferences** menu lets you do the following:

- View the user name of the user currently logged in to the system (default)
- Open the **System Settings** window, where you can configure and view license and certificate information.
- Open the **User Preferences** window, where system preferences can be set. For more information, see Customizing system preferences on page 223.
- Log out of the system
- Display information about your system, including information required for licensing (**About** option). For more information, see Viewing licensing information on page 186.

Dimmer

The **Dimmer** button toggles the dimmer feature on and off. When you use the dimmer, only tiles that are essential for real-time monitoring are lit. You can temporarily light up the dimmed tiles by hovering the mouse pointer over them.

Widget

The **Widget** button reduces the dashboard size to a widget containing a condensed display of the storage system. The widget floats on your desktop, allowing you to visually monitor your system while using other applications. The widget displays the **Capacity** tile, Workload activity, Rebuild/Rebalance activity, and Alerts indicators.

When the dashboard is minimized to a widget, the **Full-Screen** button is displayed in the top-right-corner, as shown in the following figure. This button toggles the display back to full dashboard mode.

Figure 17 Widget, showing Return to Full-Screen button

Dashboard tiles

The Dashboard tiles provide a visual overview of storage system status. The tiles are dynamic, and contents are refreshed at the interval set in the system preferences (default: 10 second intervals). System preferences can also be used to set the display to basic or advanced reporting. For more information, see Customizing system preferences on page 223.

Some of the tiles' contents differ, depending on the navigation filter in use. When the dimmer feature is enabled, the display of non-essential tiles is dimmed, unless the mouse pointer is positioned over them.

Active alert statuses relevant for specific tiles are indicated by red, orange or yellow symbols on those tiles. These indicators show that one or more alerts are active, but not the number of alerts.

Note

The filter used in the Backend view does not influence the Dashboard display.

Figure 18 Dashboard tiles

The following table describes the tiles displayed on the Dashboard.

Table 25 Dashboard tiles

Tile	Description
1 - Capacity	Displays the raw capacity of the system, rounded off to multiples of 8 GB. The available raw capacity is represented by concentric rings outwards from the center: **Outer Ring** Displays the storage usage using the color codes described in the legend. The actual values are written next to the colored segments. The icon in the bottom right corner of the tile displays the color code legend. The outer ring is divided into colors as follows: • *In use*—From the total, in clockwise direction: shows healthy, degraded and failed capacity. The total of these three items is equal to the sum of capacity in use. • *Unused*—Shows how much further the raw capacity can be expanded; if this capacity is not accessible, it will be marked as Unavailable Unused. • *Special purpose*—Shows spare capacity reserved for system operation; Decreased capacity that was deducted from devices (using the `Set Device Capacity Limit` command), and cannot be used. **Note** A miniature version of the outer Capacity ring is shown in some Backend table views and Property Sheets. **Center circle** Displays the total amount of available raw storage.

Table 25 Dashboard tiles (continued)

Tile	Description
	Note Total available raw storage does not represent the total amount of capacity available for volume allocation. **Inner Ring** Displays the snapshot usage. The arc displays the total amount of available data. The filled (bronze) part represents the capacity used by original data volumes, and the hollow (outlined) part represents the capacity used for snapshot volumes. This displays the ratio of snapshot usage. To get a more accurate idea of snapshot usage, see the Backend Capacity Usage view.
2 - I/O Workload	Displays the performance statistics of the system (IOPS, bandwidth and I/O size). More details about I/O can be viewed in the Backend table views: **Application I/O**, **Overall I/O**, and **I/O Bandwidth**. In Advanced Dashboard view (controlled by system Preferences—see Customizing system preferences on page 223) aggregated values of bandwidth and IOPS are displayed. The table in this tile summarizes the Reads, Writes and Totals of IOPS, and throughput and the average size of an I/O.
3 - Rebalance	Indicates if ScaleIO is currently redistributing data over the storage capacity. Rebalancing is usually a result of adding new storage capacity, or safely removing storage capacity, and triggers redistribution to utilize the additional compute power, and to improve performance. Rebalancing may also happen after recovery from server or storage device failure. The tile displays the rate in which the data is rebalanced, using a large green font and icon. Capacity that is pending redistribution is displayed in small white fonts.
4 - Rebuild	Indicates if ScaleIO is currently rebuilding RAID 1 data. A rebuild is usually a result of a recovery due to failure of a server or a storage device. The tile displays the rate in which the data is rebuilt, using a large orange font and icon. Capacity that is still pending rebuild is displayed in small white fonts. In Advanced Dashboard view (controlled by system User Preferences—see Customizing system preferences on page 223) more details are displayed on this tile.
5 - Alert Indicators	Displays the number of active alerts in the system, using the system-wide color codes.
6 - SDCs	Displays the number of SDCs (clients) in the system. The large number in the center is the number of SDCs connected to the MDM. The defined number includes all SDCs defined in the system (some of which may be disconnected from the MDM).
7 - Volumes	Displays the number of volumes defined across the system, the free available capacity, and the used capacity. The amount of Free capacity shown on this tile is the maximum amount that can be used for creating a new volume. This amount takes into account how much raw data is needed for maintaining RAID 1 and system spares. Note that the number of volumes and the total capacity include snapshots.
8 - Protection Domains	Displays the number and status of all Protection Domains defined in the system. The large number in the center is the number of Protection Domains. This tile is displayed when the dashboard is filtering information according to cluster.

Table 25 Dashboard tiles (continued)

Tile	Description
8 - Storage Pools	Displays the number and status of all Storage Pools defined in the Protection Domain. The large number in the center is the number of Storage Pools. This tile is displayed when the dashboard is filtering information according to Protection Domain.
9 - SDSs	Displays the number and status of all SDSs (servers) in the system. The large number in the center is the number of SDSs defined in the MDM. This tile is displayed when the dashboard is filtering according to cluster or Protection Domain. If any SDSs are currently in Maintenance Mode, the orange maintenance icon is displayed on this tile. The specific SDSs currently in Maintenance Mode can be identified using the Backend and Alerts views.
9 - Devices	Displays the number and status of all storage devices defined in the Storage Pool. The large number in the center is the number of devices defined in the MDM. This tile is displayed when the dashboard is filtering according to Storage Pool.
10 - Management	Displays the status of the MDM cluster, or of an MDM operating in Single Mode (one node). The status is displayed graphically as a combination of the MDM cluster elements, and an alert icon if active alerts exist. For more information, see Management (MDM) cluster status on page 168. When you hover your mouse pointer over this tile, a tooltip displays the IP addresses, including the Virtual IP address, used by the MDM cluster or node.
11 - Local Clock	Displays the time on your local GUI client machine (not the MDM time)

Management (MDM) cluster status

The graphics shown in the following tables represent various MDM states. A tooltip displays the IP addresses used by the MDM cluster or node.

- In 3-node cluster mode, a Master MDM, a Slave MDM, and a Tie Breaker MDM are configured in the system, and statuses are displayed on the **Management** tile of the dashboard.
- In 5-node cluster mode, a Master MDM, two Slave MDMs, and two Tie Breaker MDMs are configured in the system, and statuses are displayed on the **Management** tile of the dashboard.
- In single mode, one Master MDM is configured in the system, and the status of that MDM is displayed on the **Management** tile of the dashboard.

Note

It is not recommended to use Single Mode in production systems, except in temporary situations. The MDM contains all the metadata required for system operation. Single Mode has no protection, and exposes the system to a single point of failure. If the connection to the MDM is lost, the Dashboard is dimmed, and a dialog box is displayed.

The following figure shows how management nodes are displayed on the

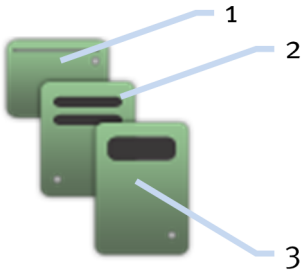

Dashboard.

Legend: 1—Tie Breaker, 2—Slave, 3—Master

The following tables show how management clusters are displayed on the Dashboard.

Table 26 Management node icons with normal operational status (green)

Icon	Description
	5-node cluster
	3-node cluster
	single node

Table 27 Management node status indications and color codes

Example	Description
	Green—Normal operation
	Gold—Degraded state: data is not consistent; system is synchronizing (the node is still operational)
	Gray—Degraded state: a Slave or Tie-Breaker is down
	Blue arrow—An upgrade is in progress
	There is no communication with the Master MDM

Frontend views

The **Frontend** view provides detailed information about frontend objects in the system, including volumes, SDCs and snapshots, and lets you perform various configuration operations. The main areas of the Frontend view are:

- **Filter**—lets you filter the information displayed in the table and Property Sheets. For more information, see Filter (1) on page 175.

- **Toolbar buttons**—let you perform commands on the selected row in the table (add, remove, configuration), or toggle the display of Property Sheets, by clicking the appropriate button
- **Table**—displays detailed information about system objects. The table displays a wide range of information, which can be filtered. Certain commands can be performed on objects, using the context-sensitive menu for the desired row in the table, or the **Command menu** on the toolbar.
- **Property Sheets**—display very detailed read-only information about the object selected in the table. For more information, see Property Sheets on page 180. For more information about the terminology used in the Property Sheets, see the Glossary at the end of this publication.

Volumes

Figure 19 Frontend > Volumes view

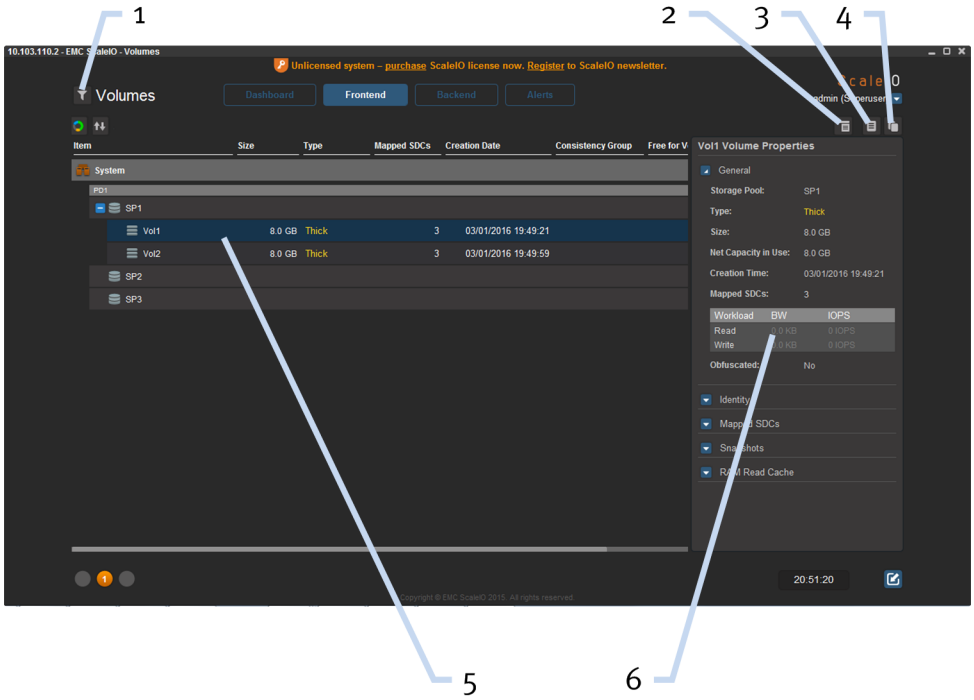

1 **Filter toggle**—controls display of the filter

2 **Table view options**—each button provides a different combination of properties which can be displayed together in the table.

3 **Command menu**—contains a list of commands which you can perform on the row selected in the table

4 **Show Property Sheet**—controls display of the Property Sheet. The Property Sheet displays information about the object selected in the table.

5 **Duplicate Property Sheet**—Opens and floats the Property Sheet in a new window, which can be kept open while you display the **Table**properties of a different object, for comparison or other purposes

6 **Property Sheet**—The Property Sheets provide detailed read-only information about the object selected in the table

SDCs

Figure 20 Frontend > SDCs view

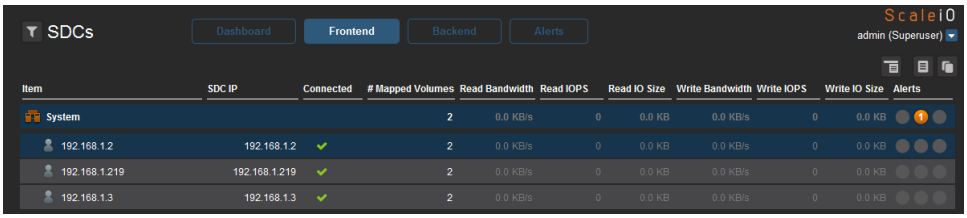

Snapshots

Figure 21 Frontend > filtered Snapshots view

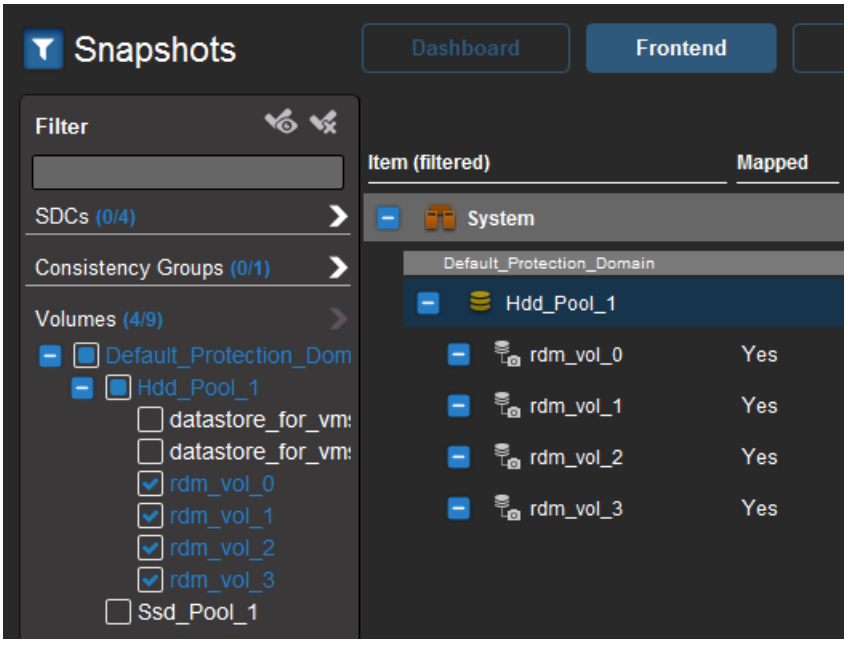

Adding Frontend objects to the filter

This topic explains how to add Frontend objects such as volumes, SDCs, snapshot trees and consistency groups to the Frontend filter. The filter simplifies the display in the main window, by showing only required objects, and hiding the rest.

Note

Toggle the filter display on and off by clicking the filter icon at the top left corner of the window. When the filter is on, only items added to the filter will be visible in the table.

To add objects to the filter, perform these steps:

Procedure

1. In the **Frontend** view, navigate to the objects that you want to add to the filter, and select them.

2. From the **Command menu** or context-sensitive menu, select one of the following, according to your needs:
   - **Add to Filter**
   - **Add Snapshot Tree to Filter** (adds all the snapshots for the selected volume)
   - **Add Consistency Group to Filter** (adds all the members of the consistency group)

Backend view

The **Backend** view provides detailed information about backend objects in the system, and lets you perform various configuration operations. The main areas of the Backend view are:

- **Filter**—lets you filter the information displayed in the table and Property Sheets. For more information, see Filter (1) on page 175.
- **Toolbar buttons**—let you display various sets of information, perform commands on the selected row in the table (add, remove, configuration), or toggle the display of Property Sheets, by clicking the appropriate button
- **Table**—displays detailed information about system objects. The table displays a wide range of information, which can be filtered. Certain commands can be performed on objects, using the context-sensitive menu for the desired row in the table, or the **Command menu** on the toolbar.
- **Property Sheets**—display very detailed read-only information about the object selected in the table. For more information, see Property Sheets on page 180. For more information about the terminology used in the Property Sheets, see the Glossary at the end of this publication.

Note

Irrelevant, mostly zero values, are “dimmed” in the Backend view. The only exception (where a value could be greater than zero but is dimmed) is when the Max Capacity is the same as the Total Capacity. In this case, Max Capacity would be dimmed (if decreased capacity exists, Max = [Total + decreased] in the Backend table). Similarly, arrows and icons are unavailable in the same circumstances.

Note

Some objects in the system can be identified by ID numbers. The ID numbers displayed in the GUI can be used in CLI commands to specify these objects.

Figure 22 Backend view

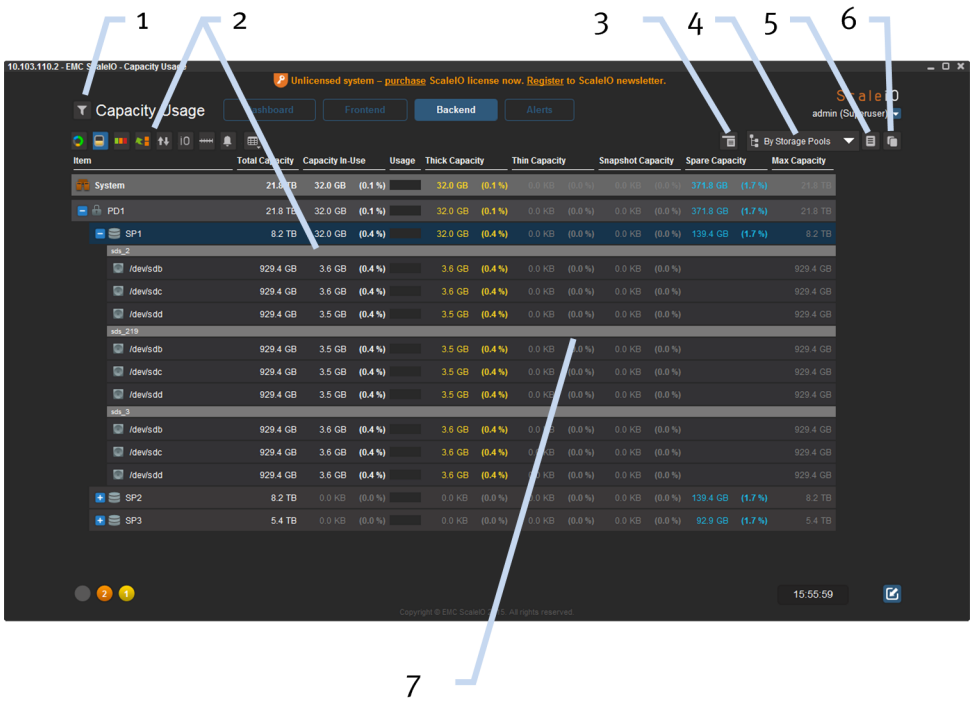

1 Filter toggle—controls display of the filter

2 Table view options—each button provides a different combination of properties which can be displayed together in the table. Additional views are available from the **More Table Views** button. For more information, see Table on page 175.

3 **Command menu**—contains a list of commands which you can perform on the row selected in the table. For more information, see Command menu (3) on page 175.

4 **Storage Pools\SDS** toggle—toggles display of the table rows grouped according to either Storage Pools, or SDSs. Some table views are only available when sorted by either SDSs or Storage Pools. For example, Fault Sets are only displayed when sorting by SDSs, and Rebuild I/O Priority is only displayed when sorting by Storage Pools.

5 **Show Property Sheet**—controls display of the Property Sheet. The Property Sheet displays information about the object selected in the table.

6 **Duplicate Property Sheet**—Opens and floats the Property Sheet in a new window, which can be kept open while you display the properties of a different object, for comparison or other purposes

7 **Table**—displays a summary of properties for the objects selected in the filter, according to the selected table view option

Command menu (3)

The **Command menu** button displays a list of commands which you can perform on rows selected in the table. The contents of the **Command menu** differ, depending on the object selected in the table. Many of the commands can also be accessed from the context-sensitive menu when table rows are right-clicked.

Filter (1)

The filter lets you filter out the information shown in the table, so that only the objects related to the items selected in the filter are visible. Items colored dark blue in the filter are displayed. The buttons at the top of the filter panel control the items hidden, displayed, selected or cleared in the filter. Below these buttons, a search field lets you type free text to search for a object by name. The filter below shows a filter being used to show specific devices in the Backend table. In the Backend, you can filter information according to Storage Pools, SDSs, Volumes, Protection Domains and Devices.

Your filtering choice would depend on your current problem or management needs. You can also add a specific object to the filter, using the `Add to Filter` command from the **Command menu** or context-sensitive menu.

Figure 23 Backend filter

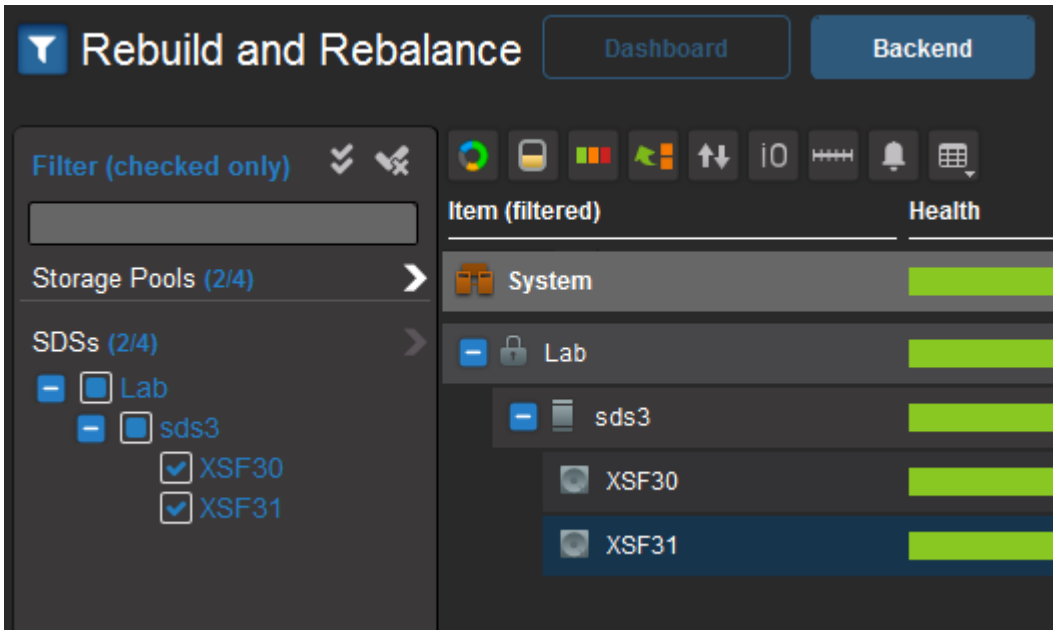

Table

The table displays object statuses, and allows configuration of some objects. The contents of a Property Sheet for a specific object are determined by the row selected in the table. The color codes described in Color codes on page 161 also apply to the items displayed in the table.

The information displayed in the table is controlled by the filter and the table view options. Selected objects (rows in the table) can be created, configured and removed using context-sensitive menu options, or commands from the **Command** menu. The columns in the table can be resized, by dragging the borders in the heading row of the

table. The scrollbar at the bottom of the table lets you scroll through the columns in the table. The following table describes the available table views:

Note

Clock icons in the cells indicate stale\aging data.

Table 28 Table view options

Name of view	Contents of view	Suggested use, comments
Overview	Total Capacity, Capacity In-Use, I/O Bandwidth, IOPS, Rebuild, Rebalance, Alerts	Provides a general overview the capacity and health of system objects. If you find that there are active alerts, you can switch to **State Summary** table view, or the **Alerts** view to see more details about the alerts. While this view is similar to the Dashboard, the main difference is that here, you can simultaneously monitor many objects within their hierarchy. You can then filter the table for specific objects or sets of objects, in contrast to the Dashboard, where you can only drill-down to objects. In addition, commands and Property Sheets are available to you from this view. **Note** The miniature Capacity ring shown here represents the outer Capacity ring on the Dashboard. For more information, see Dashboard tiles on page 165.
Capacity Usage	Total, In-Use, Usage, Thick Thin, Snapshot, Spare, Max Related Property Sheet section: **Capacity**	Provides a breakdown of capacity usage per use type. You can use this to check whether more capacity needs to be added to your system, and where.
Capacity Health	Capacity In-Use, Protected, Degraded, Failed, Health, Rebuild, Rebalance Related Property Sheet sections: **Capacity**, **Alerts**, **Rebuild/Rebalance**	Provides information about the health of the objects in the system, per object
Rebuild and Rebalance	Health, Backward Rebuild, Forward Rebuild, Rebalance Related Property Sheet sections: **Alerts**, **Rebuild/Rebalance**	Provides a summary of Rebuild and Rebalance health, status, and workload per object
Application I/O	Bandwidth, IOPS and I/O Size for: Total, Read, Write, 2nd Write Related Property Sheet section: **Workload**	Provides workload information for applications reading/writing to storage in the system. 2nd Writes refer to the protection copy of data being written to storage.
Overall I/O	Bandwidth, IOPS and I/O Size for: Total, Total Read, Total Write Related Property Sheet section: **Workload**	Provides workload information for all I/Os in the system, including both application I/Os, and I/Os for internal processes

Table 28 Table view options (continued)

Name of view	Contents of view	Suggested use, comments
I/O Bandwidth	Read, Write, Backward Rebuild, Forward Rebuild, Rebalance, Total, Total Read, Total Write Related Property Sheet section: **Network Throttling**	Shows the bandwidth being used for various jobs in the system
State Summary	Summary Related Property Sheet section: **Alerts**	Can be used to identify items in the system which have open alert states, and to view the alert messages, including information statuses of various objects. If there are pending security certificates, and if an SDS is in Maintenance Mode, this is also indicated here. Alerts marked blue are for information purposes only, and do not require that any action be taken.
Configuration	Total Capacity, SDSs, Devices, Storage Pools, Volumes, Free Capacity for Volume Allocation, Alerts Related Property Sheet sections: **Identity**, **Related Objects**	Provides an overview of the number of objects per type in your system, their capacity, and lets you determine the amount of free capacity available for creating an additional volume. The **Free Capacity for Volume Allocation** column is the only one that is not stated in Raw Capacity values. This table view is a convenient location from which to perform Add, Remove, Activate and Inactivate commands. For more information, see Configuring the System using the GUI on page 187. The **Related Objects** section of the Property Sheet helps you to identify related objects to the one selected in the table.
Device Details	Average Read Latency, Average Write latency, Average Read Size, Average Write Size, Scanned Capacity, Resolved Errors, Data Conflicts Related Property Sheet sections: **Device Latency** , **Background Device Scanner**	Provides an overview of performance information, and Background Device Scanner results, per device. When the table is sorted by SDSs, the Down arrow in each SDS row reveals all the devices in the SDS. **Note** Information is shown only at Device level, and there is no aggregation of information at higher levels. If the background device scanner is enabled, several device read statistics are dramatically affected.
RAM Read Cache	State, Size, Used, Hit Rate, Write Mode Related Property Sheet section: **RAM Read Cache**	Lets you check which SDSs have RAM Read Cache enabled, view associated counters, and check which Storage Pool is set to use the cache for its devices. Advanced feature; modify with caution.

Table 28 Table view options (continued)

Name of view	Contents of view	Suggested use, comments
Read Flash Cache	State, Size, Used, Read Hit Rate, Total Errors, Alerts Related Property Sheet sections: **Read Flash Cache** and **Related Objects**	Lets you check which SDSs and Storage Pools have Read Flash Cache enabled, view associated counters. The **Related Objects** section of the corresponding Property Sheet shows the names/paths of the devices used for caching.
Rebuild and Rebalance (Detailed)	Health, Backward Rebuild, Forward Rebuild, Rebalance Related Property Sheet section: **Rebuild/Rebalance**	Shows workload and active rebuild and rebalance jobs, as well as pending jobs to be processed. Credited incoming and outgoing information is displayed for each process (Backward Rebuild, Forward Rebuild and Rebalance).
Planned Rebuilds (Advanced)	Degraded, Distribution, Backward and Forward Rebuilds Related Property Sheet section: **Rebuild/Rebalance**	Shows amount of unprotected data, a visual breakdown of rebuild progress (Distribution), bandwidth, direction and status (Active/Pending) of jobs. You can monitor how the degraded capacity is planned to be rebuilt: How failed degraded capacity is expected to be rebuilt, then how the system will pool it for the rebuild processes (backward and forward,) and which of it is actively being rebuilt. **Note** Overall Degraded capacity includes both healthy and failed copies of the data, while the processed capacity includes only the failed copies. Advanced feature; modify with caution.
Planned Rebalancing (Advanced)	Protected, Distribution, Rebalance Related Property Sheet section: **Rebuild/Rebalance**	Shows amount of protected data, a visual breakdown of rebalance progress (Distribution), bandwidth, direction and status (Active/Pending) of jobs. Advanced feature; modify with caution.
Rebuild I/O Priority (Advanced)	Policy, Concurrent I/O limit, Bandwidth Limit, Application Threshold, Quiet Period Related Property Sheet section: **I/O Priority**	Displays bandwidth settings currently configured for Rebuild jobs. These I/O Priority settings apply only to Storage Pools. These settings control system performance. Advanced feature; modify with caution.
Rebalance I/O Priority (Advanced)	Policy, Concurrent I/O limit, Bandwidth Limit, Application Threshold, Quiet Period Related Property Sheet section: **I/O Priority**	Displays I/O priority settings currently configured for Rebalance jobs. These I/O Priority settings apply only to Storage Pools. These settings control system performance. Advanced feature; modify with caution.
Network Throttling (Advanced)	Overall I/O Limit, Rebuild I/O Limit, Rebalance I/O Limit, Rebuild Incoming Limit, Rebalance Incoming Limit, Rebuild Queue Length, Rebalance Queue Length, Rebuild Outgoing Jobs, Rebalance Outgoing Jobs	Displays Network Throttling settings currently configured in the system. These settings control system performance. Advanced feature; modify with caution.

Table 28 Table view options (continued)

Name of view	Contents of view	Suggested use, comments
	Related Property Sheet section: **Network Throttling**	
RAM Read Cache (Internal)		For Support purposes only. Visible only if Backend Internals are enabled in the **User Preferences** window.

Alerts view

The **Alerts** view provides a list of the alert messages currently active in the system, in table format. You can filter the table rows according to alert severity, and according to object types in the system. For a list of alerts generated by the system, see ScaleIO Alerts in SNMP, GUI, REST, and ESRS on page 535

To view a Property Sheet for a specific alert, select the corresponding row in the table, and click the **Show Property Sheet** button. For more information about Property Sheets, see Property Sheets on page 180.

Figure 24 Alerts view

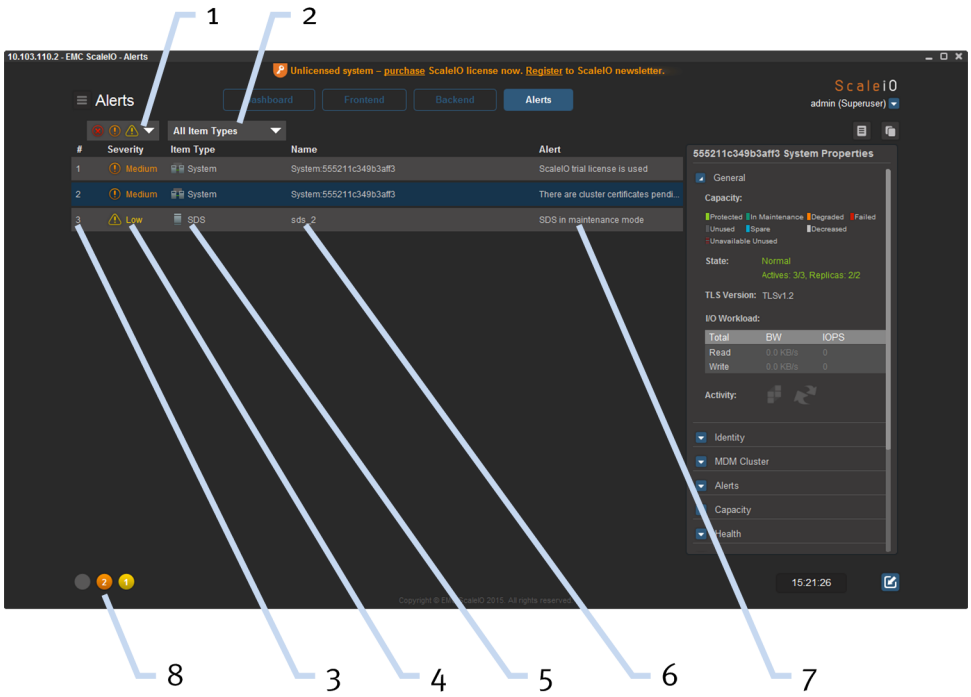

1 Severity filter—filters the table contents according to alert severity:

All alerts

Medium—alerts which are Medium or High severity

5 **Item Type**—MDM, Protection Domain, Storage Pool, SDS, Device, SDC

	High—only alerts which are High severity		
2	Item Types filter—filters the table contents according to: All Item Types, MDM, Protection Domain, Storage Pool, SDS, Device, SDC	6	**Name**—the user-defined name of the item, if one has been defined
3	**#**—the line number in the Alerts table	7	**Alert**—the alert message generated by the system
4	**Severity**—the alert severity: High, Medium or Low	8	Alert indicators—summarize the total amount of each alert type (from left to right: High, Medium, Low)

Property Sheets

The Property Sheets provide detailed read-only information about the object selected in the Frontend table, the Backend table, or in the Alerts table.

The contents of the Property Sheets differ, depending on the object selected in the table. Property sheets help you to monitor specific objects in the system by displaying the following, using the blue collapse and expand arrows next to each section of the Property Sheet:

- **General** and **Health** information about the object
- **Identity**—identifying information, such as an object's ID number, name, IP addresses, port usage, VM usage, GUID
- **Mapped SDCs** for the selected volume, including the SDC name/IP address, and bandwidth limit and IOPS limit per volume
- **Mapped Volumes** for the selected SDC, including the volume name, and bandwidth limit and IOPS limit per volume
- **Snapshots**—details about net capacity in use, and a snapshot tree including the volume name, and date and time of creation of the parent volume and each snapshot
- **MDM Cluster** details, such as cluster mode and state, IP addresses and ports (including the Virtual IP address) used for management and MDM functions, and SSL version being used for secure communication. The Virtual IP addresses and the Virtual IP interfaces are displayed together in table format, to indicate the mapping between them.
- **Alerts**—alert status per selected object, including SDS Maintenance Mode (if active)
- **Capacity** usage per selected object
- **Workload** information (bandwidth usage and IOPS) per selected object
- **Rebuild/Rebalance** information about the selected object(s), for forward and backward rebuilds, rebalancing, data at rest, and job status (active\pending). The type of information displayed depends on the type of object selected in the table.
- **Device Latency** averages for Read and Write, and average I/O size, for the device selected in the table

Note

If the background device scanner is enabled, several device read statistics are dramatically affected.

- **Consistency Group** for the related volumes, including a list of the volumes, and the Consistency Group name
- **RAM Read Cache** configuration, state, and statistics for the selected object. The type of information displayed depends on the type of object selected in the table.
- **Read Flash Cache** configuration, state, and statistics for the selected object. The type of information displayed depends on the type of object selected in the table.
- **Device Test Results** for the device selected in the table, if any tests have been performed
- **Background Device Scanner** results for the device selected in the table, if the scanner is enabled
- **Network throttling** configuration for the selected SDS, including bandwidth limit per job, and queue length, for both Rebuilding and Rebalancing.
- **I/O Priority** configured for the selected Storage Pool, including Rebuild and Rebalance states, number of parallel jobs, I/O prioritization policy, concurrent I/Os and bandwidth limit
- Miscellaneous items, such as DRL mode, zero padding, checksum mode
- **Performance** profile currently assigned to the selected object: Default, High or Custom. “Custom” is displayed whenever performance-related parameters have been configured manually, instead of via a performance profile.
- **Oscillating Failure Counters** are shown for SDSs and devices selected in the Backend view, for SDCs selected in the Frontend view, and for Alert messages.
- **Oscillating Failure Parameters** currently configured in the system are shown for the entire system, for the selected Storage Pool or Protection Domain, and for Alert messages. The counters shown depend on the object selected in the table.
  - **Window:** the sliding time-window for each interval (Short, Medium and Long)
  - **Threshold:** the number of errors that may occur before error reporting commences
  - **Period:** the time interval of each Window, in seconds
- **Maintenance Mode** state of the selected SDS is shown here. States include: No Maintenance and In Maintenance Mode.
- **Related objects**, which can be very useful for troubleshooting problems, or for planning purposes, when you need to make changes to your system.

You can view properties for multiple objects by using the **Duplicate Property Sheet** button, and then navigating to a different object’s row in the table. When more than one Property Sheet is open, a floating widget that controls them is displayed in the bottom part of the main window, as shown in the following figure.

Note

You can open Property Sheets, duplicate them, and then simultaneously work on other unrelated objects in the system. The duplication feature is not supported in the Hardware view.

Figure 25 Multiple floating Property Sheets

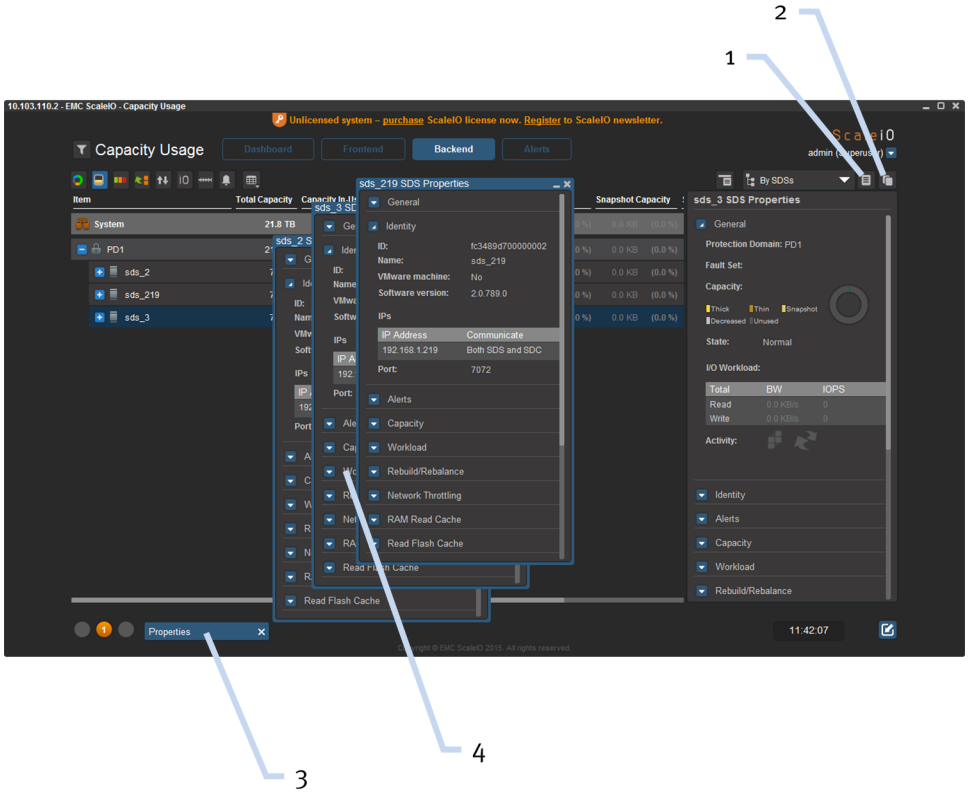

1 **Display\Hide Property Sheet** button

2 **Duplicate Property Sheet** button

3 **Duplicated Property Sheets** toggle

4 Multiple Property Sheets opened using the **Duplicate Property Sheet** button

The following figure shows a typical Property Sheet for an SDS, with the **Related Objects** section of the Property Sheet expanded.

Figure 26 Example of a Property Sheet for an SDS

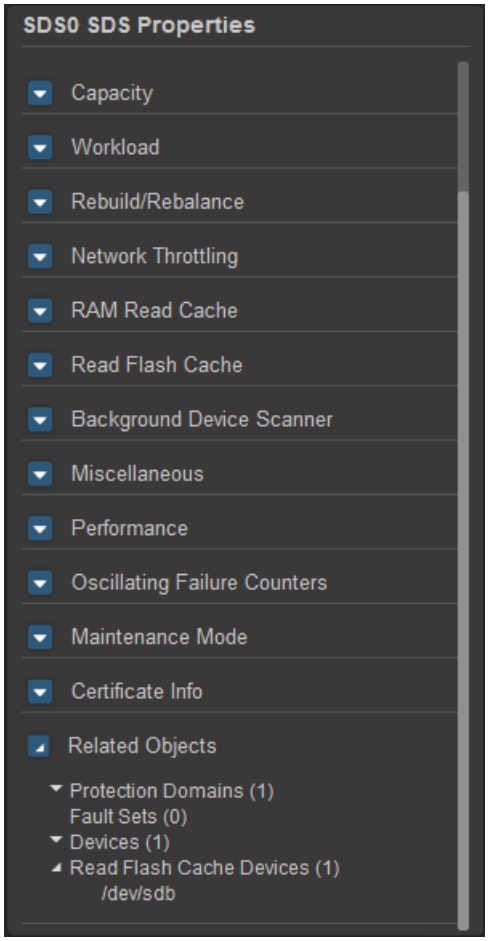

CHAPTER 11

Monitoring the System using the GUI

This chapter explains how to monitor the system, using the GUI. Topics include:

- Viewing object properties in Backend and Frontend views.....186
- Viewing licensing information..... 186
- Verifying your connection to the Management cluster..... 186

Viewing object properties in Backend and Frontend views

Procedure

1. In **Backend** or **Frontend** view, use the filter to display one or more objects, and select the corresponding check boxes of the desired objects (optional).
2. Navigate to the desired object in the table.
3. Display the required information, using the table view options.
4. Select the required object's row in the table, and then, on the expandable Property Sheet on the right side of the window, click the blue arrow buttons beside the headings to expand them and view specific status information.

Note

Contents of the Property Sheet are dynamic, and differ, depending on the row selected in the table. For more information about Property Sheets, see Property Sheets on page 180.

Viewing licensing information

Before a permanent license is installed, a banner is displayed at the top of the GUI window that provides links for purchasing a license and for registering for ScaleIO newsletters.

Information required for licensing purposes is located in the **About** window, as described in the following steps:

Procedure

1. From the **Preferences** menu at the top right side of the window, in any view, select the **About** option.

   The **About** window is displayed.
2. Make a note of the information displayed for **Installation ID**, which is required for electronic licensing purposes.

   Additional information pertaining to your license is also displayed in this window.

Verifying your connection to the Management cluster

If your GUI connection to the system is operating normally, regular access to all the GUI views is possible.

The Management IP address to which you are connected is also displayed in the top left corner of the GUI. For more information about MDM statuses, see Management (MDM) cluster status on page 168.

If the connection is lost, the GUI is dimmed, and a dialog box is displayed.

CHAPTER 12

Configuring the System using the GUI

This chapter explains how to configure the system, using the GUI. Topics include:

- Configuring capacity 188
- Configuring cache 195
- Configuring volumes, volume trees, SDCs, and snapshots 200
- Entering and exiting SDS Instant Maintenance Mode 210
- Configuring and viewing Oscillating Failure counters 211
- Applying Performance Profiles to system components 216
- Configuring I/O priorities and bandwidth use (advanced) 217
- Enabling and disabling Rebuild\Rebalance (advanced) 219
- Using the background device scanner 219
- Modifying Checksum protection mode 221
- Renaming objects 222
- Approving pending security certificates 223
- Customizing system preferences 223

Configuring capacity

Add, remove, and configure capacity.

The following topics explain how to add, remove, activate, and inactivate capacity, activate devices, clear device errors, and set device capacity limits.

Adding SDSs and storage devices

SDSs and storage devices can be added to a system one by one, or in bulk operations, using the `Add SDS` and `Add Device` commands. In addition, the **Add SDS to Protection Domain** window lets you add both SDSs, and corresponding devices, all from the same window. You can associate up to eight IP addresses to the SDS. By default, performance tests are performed on the added devices, and the results are saved in the system.

You can assign a name to the SDS, as well as to the devices. This name can assist in future object identification. This can be particularly helpful for SDS devices, as the name will remain constant, even if the path changes.

To run CLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *ScaleIO CLI Reference Guide*.

Before you begin, ensure that at least one suitable Storage Pool is defined in the required Protection Domain.

Note

Devices can be tested before going online. Various testing options are available the **Advanced** area of the window (default: **Test and Activate**).

Note

You cannot enable zero padding after adding the devices. For more information, see Storage Pools on page 37, as well as the Command Quick Reference on page 99.

Procedure

1. In the **Backend** view:
   a. To add one or more SDSs, navigate to the required Protection Domain, and select its row in the table
   b. To add one or more storage devices to an existing SDS, navigate to the required SDS, and select its row in the table
2. From the **Command menu** or context-sensitive menu, select the desired **Add** option.

   An **Add** window is displayed.
3. Enter the relevant information in the fields. Object names must meet the following requirements:
   - Contains less than 32 characters
   - Contains only alphanumeric and punctuation characters
   - Is unique within the object type

a. Fields that contain orange explanation marks are mandatory.

b. You must add at least one device to the new SDS at this stage.

   You can add more devices later.

c. If you add Read Flash Cache devices, ensure that caching policy is set to enabled in the corresponding Storage Pool and on the SDS.

   For more information, see Setting Read Flash Cache policy at Storage Pool level on page 195 and Setting Read Flash Cache policy at SDS level on page 196.

   Note

   If you want to add an SDS without any devices, you can do so using the CLI. To run CLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *ScaleIO CLI Reference Guide*.

d. The **Advanced** option provides additional items, such as device testing and RAM Read Cache configuration.

   Click its **Expand** button to display additional fields, and configure them (recommended for advanced users only).

e. For some object types, a button is displayed.

   Click it to add more objects or rows of the same type.

4. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

5. Click **Close**.

   If you chose the **Test only** option in step 3, activate the devices as described in Activating devices on page 194.

Figure 27 Add SDS window

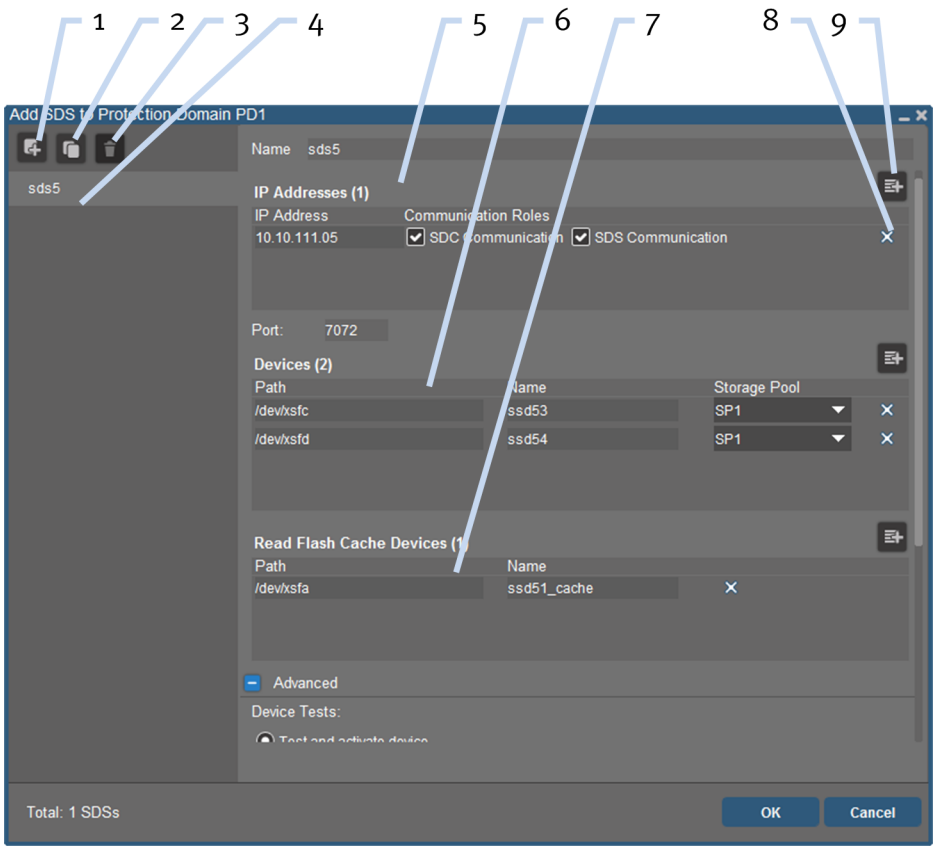

1	Add another SDS	6	Storage devices and their properties
2	Add another SDS by duplicating the values in the right panel of the window	7	Read Flash Cache devices and their properties
3	Remove the selected SDS from the window	8	Add object/row
4	List of SDSs that will be added to the Protection Domain	9	Remove object/row
5	SDS properties		

Figure 28 Add Device window showing command validation

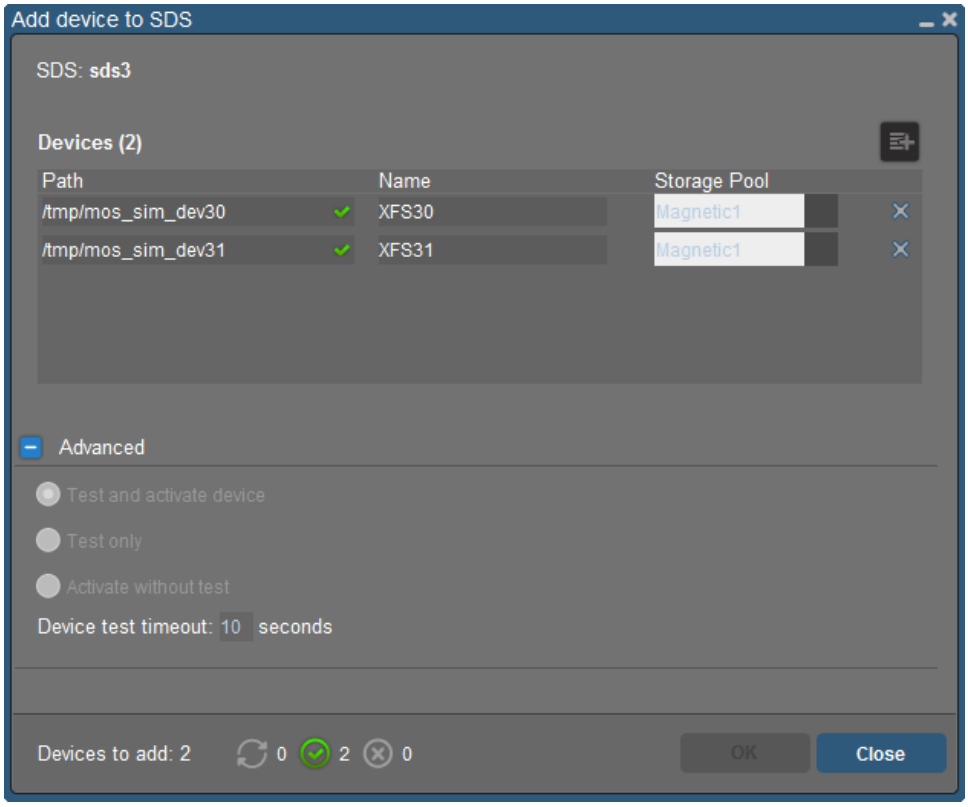

Removing SDSs and devices

The removal of some objects in the system can take a long time, because removal may require data to be moved to other storage devices in the system. If you plan to replace a device with a device containing less storage capacity, you can configure the device to a smaller capacity than its actual capacity, in preparation for replacement. This will reduce rebuilding and rebalancing operations in the system later on. For more information, see Setting device capacity limits on page 195.

The system has job queues for operations that take a long time to execute. You can view the jobs in the **Planned Rebuilds** and **Planned Rebalancing** table views. Operations that are waiting in the job queue are shown as Pending. If a job in the queue will take a long time, and you do not want to wait, you can cancel the operation using the **Abort** button in the **Remove** command window (if you left it open), or using the `Abort` command from the **Command menu**.

The `Remove` command deletes the specified objects from the system. Use the `Remove` command with caution.

Procedure

1. In the **Backend** view, navigate to the desired object in the table, and select its row.
2. From the **Command menu** or context-sensitive menu, select the desired `Remove` command.

   In the confirmation window, click **OK**.The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation. For some objects, an **Abort** button is available in the window, which can be used

if you decide to abort the operation. The is also an `Abort` command accessible from the **Command menu**.

3. Click **Close**.

Adding, removing, and activating and inactivating capacity

Extra storage capacity can be added to your ScaleIO system by adding SDSs and/or their storage devices. You can either add them to existing Protection Domains and Storage Pools, or create new ones.

The Dashboard **Capacity** tile, some Backend table views (such as **Capacity Usage**, **Configuration**), and Property Sheets help you to better understand the amount of raw capacity and net free capacity currently available in the system.

Adding, removing, activating and inactivating Protection Domains

This section describes how to add, remove, activate and inactivate Protection Domains. Inactivating a Protection Domain does not remove it from the system, but it makes all data stored in that Protection Domain inaccessible to the system.

The inactivation feature is a much more effective way to shut down nodes, and is preferable to shutting them down manually.

When this feature is in effect, the following activities can take place, behind the scenes:

- Determine if there are any current rebuild/rebalance activities taking place. If so, the shutdown will be delayed (unless it is forced) until they are finished.
- Block future rebuild/rebalance activities.
- Quiesce (temporarily disable) application I/O and disable access to volumes.
- Move the DRL mode of all SDSs to harden, in preparation for rebooting the server.
- Reload of all SDSs before re-enabling data access.

For each of the following procedures, after you click **OK**, the progress and result of the operation is displayed at the bottom of the window.

Note

Protection Domain removal is only possible if the Protection Domain is empty. If you inactivate a Protection Domain, the data remains on the SDSs, and therefore, it is preferable to remove a Protection Domain if you no longer need it.

To add a Protection Domain, perform these steps:

Procedure

1. In the **Backend** view, select the **System** row.
2. From the **Command menu** or context-sensitive menu, select **Add Protection Domain**.

   The **Add Protection Domain** window is displayed.
3. Type a name in the **Name** box, and click **OK**.

   When the operation is complete, the Protection Domain is active. You can now add SDSs, Storage Pools and Fault Sets to the Protection Domain. Before you add devices, ensure that at least one suitable Storage Pool is defined in the Protection Domain. To remove Protection Domains, perform these steps:

4. In the **Backend** view, navigate to, and select one or more Protection Domains.
5. Verify that you have removed all child nodes from the Protection Domain.
6. From the **Command menu** or context-sensitive menu, select **Remove.**
7. Click **OK**.
8. If a confirmation window appears, confirm the operation, and type your password if requested to do so.

   To inactivate Protection Domains, perform these steps:
9. In the **Backend** view, navigate to, and select one or more Protection Domains.
10. From the **Command menu** or context-sensitive menu, select **Inactivate Protection Domain.**

    The **Inactivate Protection Domain** window is displayed.
11. Click **OK**.

    If a confirmation window appears, confirm the operation, and type your password if requested to do so.To activate Protection Domains, perform these steps:
12. In the **Backend** view, navigate to, and select one or more Protection Domains.
13. From the **Command menu** or context-sensitive menu, select **Activate Protection Domain.**

    The **Activate Protection Domain** window is displayed.
14. Click **OK**.

Adding Fault Sets

Fault Sets provide additional safeguards for protecting your data against hardware failure. Fault Sets are subsets of a given Protection Domain.

Note

When defining Fault Sets, you must follow the guidelines described in Fault Sets on page 38. Failure to do so may prevent creation of volumes.

Procedure

1. In the **Backend** view, navigate to, and select the Protection Domain.
2. From the **Command menu** or context-sensitive menu, select **Add Fault Set**.
3. Type a name in the **Fault Set Name** box, and click **OK**.

   The Fault Set will now be visible in the **Related Objects** section of the Protection Domain's Property Sheet.

   Note

   Use the CLI to remove Fault Sets. For more information, see the Command Quick Reference on page 99.

Adding and removing Storage Pools

A Storage Pool is a group of devices within a Protection Domain. Each time that you add devices to the system, you must map them to Storage Pools. Create Storage Pools before you start adding devices to the system.

Note

You cannot enable zero padding after adding the devices. For more information, see Storage Pools on page 37, as well as the Command Quick Reference on page 99.

To add a Storage Pool, perform these steps:

Procedure

1. In the **Backend** view, select the desired Protection Domain.
2. From the **Command menu** or context-sensitive menu, select **Add Storage Pool**.

   The **Add Storage Pool** window is displayed.
3. Type a name in the **Name** box.

   Note

   The RAM Read Cache features are advanced features, and it is usually recommended to accept the default values. You can configure these features later, if necessary, using the `Configure RAM Read Cache` command. For more information about RAM Read Cache features, see Managing RAM read cache on page 134
4. Click **OK**.

   To remove a Storage Pool, perform these steps:
5. In the **Backend** view, select the Storage Pool.
6. From the **Command menu** or context-sensitive menu, select **Remove**.

   The **Remove Storage Pool** window is displayed.
7. Click **OK**.

Activating devices

Use the `Activate Device` command in the following situations:

- Devices were added to the system using the **Test only** option for Device Tests, and successfully passed the tests.
- Devices were inactivated, and you want to bring them back online.

Procedure

1. In the **Backend** view, navigate to the device or devices in the table, and select the corresponding rows.
2. From the **Command menu** or context-sensitive menu, select **Activate Device**.

Clearing device errors

Procedure

1. In the **Backend** view, navigate to the device in the table, and select its row.
2. From the **Command menu** or context-sensitive menu, select **Clear Device Errors**.

Setting device capacity limits

In circumstances where you need to replace a device in your system with a device of a smaller capacity, you should first set the capacity limit of the device to be removed to less than its full capacity. In such a case, capacity will be decreased, but the size of the disk remains unchanged. The capacity assigned to the SDS device must be smaller than its actual physical size.

Note

Decreased capacity is shown on the Dashboard, using pale gray, on the outer ring on the Capacity tile.

Procedure

1. In the **Backend** view, navigate to the device in the table, and select its row.
2. From the **Command menu** or context-sensitive menu, select **Set Device Capacity Limit**.
3. Type the desired value and click **OK**.

Configuring cache

Configure the system's caching features to optimize system performance.

The following procedures explain how to configure the system's caching features:

- Setting Read Flash Cache policy at Storage Pool level on page 195
- Setting Read Flash Cache policy at SDS level on page 196
- Adding Read Flash Cache devices on page 197
- Removing Read Flash Cache devices on page 198
- Changing RAM Read Cache volume settings on page 197
- Configuring RAM Read Cache (advanced, Backend) on page 198

Setting Read Flash Cache policy at Storage Pool level

This topic describes how to enable and disable Read Flash Cache policy at Storage Pool level.

- Once this is enabled at Storage Pool level, set the Read Flash Cache policy at SDS level.
- Once Read Flash Cache is enabled at both Storage Pool and SDS levels, add Read Flash Cache devices to the SDS in order to commence caching.

  To set Read Flash Cache policy at Storage Pool level, perform these steps:

Procedure

1. In the **Backend** view, navigate to the required Storage Pool, and display the **Read Flash Cache** table view.

The **State** column indicates whether the Read Flash Cache feature is in use in the Storage Pool.

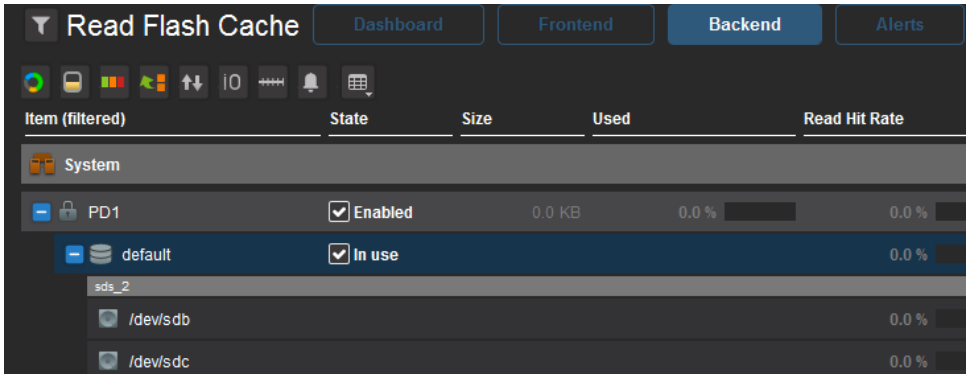

2. From the **Backend** view, right-click the Storage Pool, and choose **Set Flash Cache Policy**.
3. In the **Set Flash Cache Policy** window, do one of the following:
   a. To enable caching, ensure that the **Enable Read Flash Cache** check box is selected, and click **OK**.
   b. To disable caching, clear the **Enable Read Flash Cache** check box, and click **OK**.
4. When the status shows that the operation was successful, click **Close**.

Setting Read Flash Cache policy at SDS level

This topic describes how to enable and disable Read Flash Cache policy on an SDS.

- To enable caching, ensure that the policy is also enabled at Storage Pool level.
- Once Read Flash Cache is enabled at both Storage Pool and SDS levels, add Read Flash Cache devices to the SDS in order to commence caching.

To set Read Flash Cache policy on an SDS, perform these steps:

Procedure

1. In the **Backend** view, navigate to the required SDS, and display the **Read Flash Cache** table view.

   The **State** column indicates whether the Read Flash Cache feature is in use on the SDS:

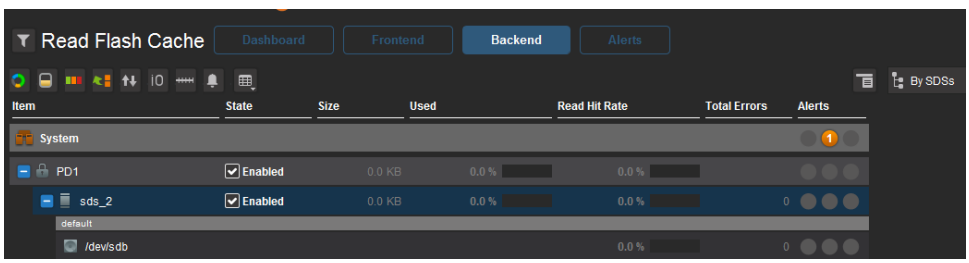

   .

2. From the **Backend** view, right-click the SDS, and choose **Set Flash Cache Policy**.
3. In the **Set Flash Cache Policy** window, do one of the following:
   a. To enable caching, ensure that the **Enable Read Flash Cache** checkbox is selected, and click **OK**.

b. To disable caching, clear the **Enable Read Flash Cache** checkbox, and click **OK**.

4. When the status shows that the operation was successful, click **Close**.

Adding Read Flash Cache devices

Up to eight caching devices can be used per SDS. Ensure that Read Flash Cache policy is set to enabled on the corresponding Storage Pool and SDS.

Note

The RFcache driver must be installed before you can add a Read Flash Cache device to an SDS. (Typically, the driver is installed during deployment, and devices are designated for caching.)

To add a Read Flash Cache device to the system, perform these steps:

Procedure

1. From the **Backend** view, navigate to the corresponding SDS, right-click it and choose **Add Read Flash Cache Device**.
2. In the **Add Read Flash Cache Device to SDS** window, type the path to the required device in the **Path** box, and a name in the **Name** box (optional).

   a. If you want to add multiple cache devices, click **Add Device** and repeat this step.

3. Click **OK**.
4. When the status shows that the operation was successful, click **Close**.
5. In some cases, a rebuild/rebalance begins.

   Use the **Dashboard** view to determine when that is complete. The cache device has been added to the ScaleIO system.

   Note

   Devices used for caching are not shown in the Backend tables. You can identify them in the **Related Objects** section of an SDS's Property Sheet.

Changing RAM Read Cache volume settings

By default, RAM Read Cache is disabled on volumes. To change RAM Read Cache settings on volumes, perform these steps:

Procedure

1. In the **Frontend > Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Set Volume RAM Read Cache**.

   The **Set Volume RAM Read Cache** window is displayed, showing a list of the volumes that will be modified.

3. Select or clear the **Use RAM Read Cache** check box as follows:

   a. To disable RAM Read Cache on the volumes, clear the check box.

b. To enable RAM Read Cache on the volumes, select the check box.

4. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

   Figure 29 Set Volume Use RAM Read Cache window

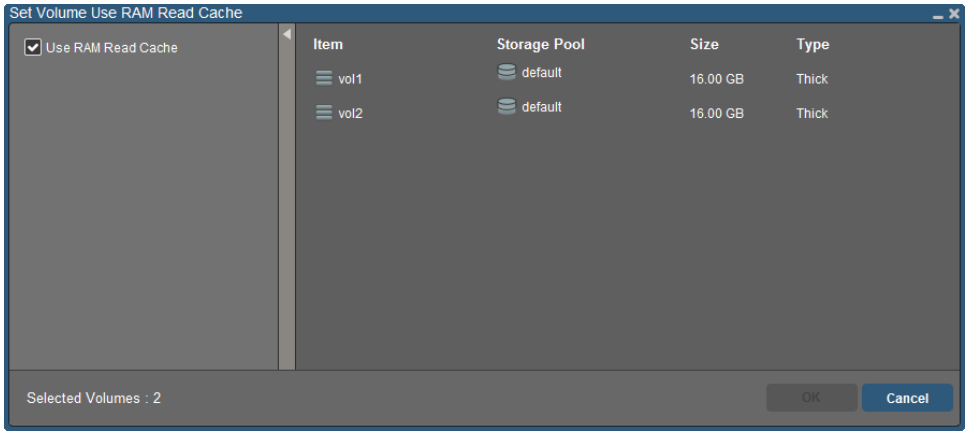

Removing Read Flash Cache devices

This topic describes how to remove an SSD device or a PCIe flash disk that is being used to provide caching.

To remove a Read Flash Cache device from the system, perform these steps:

Procedure

1. From the **Backend** view, right-click the Storage Pool of which the disk is a member, and choose **Set Flash Cache Policy**.
2. In the **Set Flash Cache Policy** window, clear the **Enable Read Flash Cache** check box, and click **OK**.
3. When the status shows that the operation was successful, click **Close**.
4. From the **Backend** view, navigate to the corresponding SDS, right-click it and choose **Remove Read Flash Cache Device**.
5. In the **Remove Read Flash Cache Device** window, click **OK**.
6. When the status shows that the operation was successful, click **Close**.
7. In some cases, a rebuild/rebalance begins.

   Use the **Dashboard** view to determine when that is complete. The device is removed from the ScaleIO system. You may now either remove the physical drive from the system, or add it to the SDS as a storage device. If you have other cache devices installed in the SDS, set the Read Flash Cache policy back to Enabled.

Configuring RAM Read Cache (advanced, Backend)

The RAM Read Cache feature improves your system's application performance for storage-related activities. By default, caching is disabled.

To use RAM Read Cache, you need to configure RAM Read Cache settings at two levels:

- Storage Pool—controls RAM Read Cache for all the SDSs in the selected Storage Pool. Caching can be enabled or disabled, and either Cached (default) or Passthrough Write Handling modes can be selected. When RAM Read Cache is enabled in a Storage Pool, the feature is enabled at Storage Pool level. However, caching must also be set to Enabled in each SDS in the Storage Pool. Caching will only begin once devices have been added to the SDSs. It is possible to enable RAM caching for a Storage Pool and then disable caching on one or more SDSs individually.
- Per SDS—controls RAM Read Cache for one or more SDSs. Caching can be enabled or disabled for the specified SDS, and the capacity allocated for caching on an SDS can be specified. Caching will only begin after one or more devices are added to the SDSs. Ensure that the feature is also enabled at Storage Pool level.

Note

By default, RAM read cache is disabled in all volumes. You can change this setting using the CLI. For more information, see Command Quick Reference on page 99.

To configure caching, perform these steps:

Procedure

1. In the **Backend** view, navigate to, and select the desired Storage Pools.
2. From the **Command menu** or context-sensitive menu, select **Configure RAM Read Cache**.

   The **Configure RAM Read Cache** window is displayed. The right pane of the window lists the Storage Pools that you are configuring.
3. Select or clear the options that you require (selected=used; clear=not used), and click **OK**.

   Figure 30 RAM Read Cache configuration at Storage Pool level

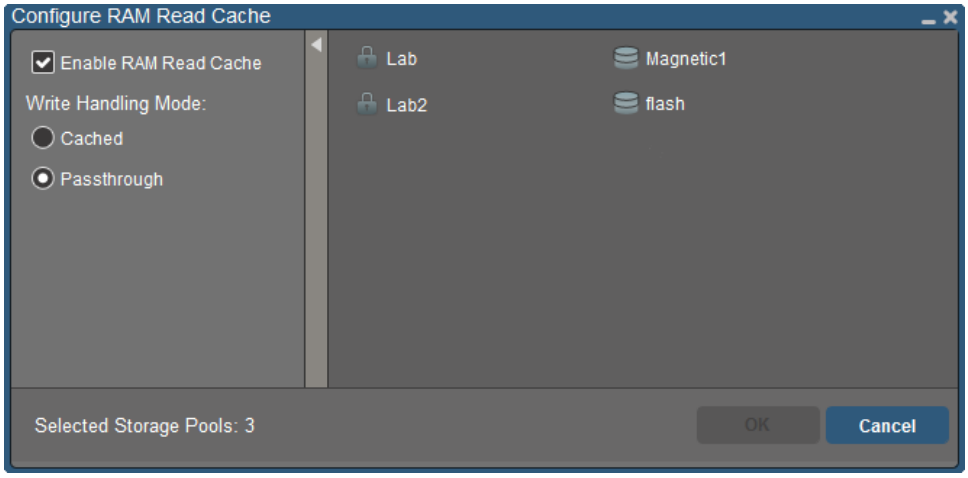

4. To enable\disable\configure cache size for SDSs, in the **Backend** view, navigate to, and select the desired SDS(s).
5. From the **Command menu** or context-sensitive menu, select **Configure RAM Read Cache**.

   The **Configure RAM Read Cache** window is displayed. The right pane of the window lists the SDSs that you are configuring.
6. Select or clear the option that you require (selected=enable; clear=disable).

7. If necessary, edit the value in the **RAM Read Cache Size** box (default=128 MB).
8. Click **OK**.

Figure 31 RAM Read Cache configuration at SDS level

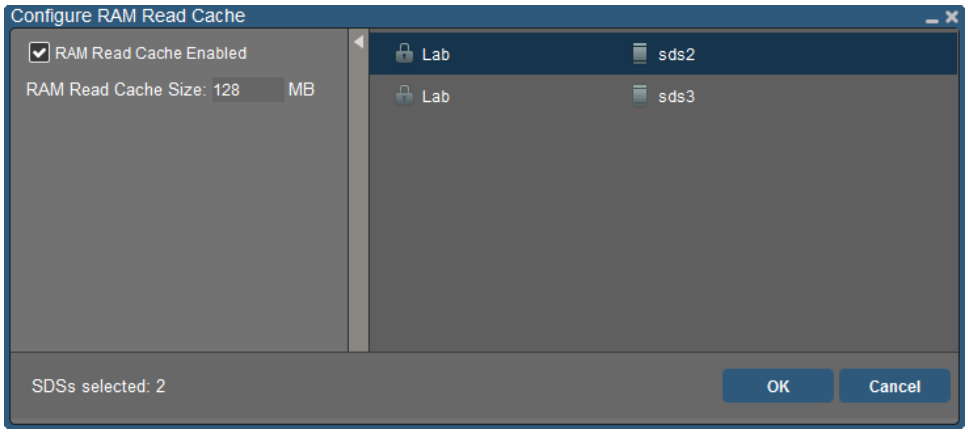

Configuring volumes, volume trees, SDCs, and snapshots

This section contains procedures for adding, removing, and managing volumes and snapshots. It also explains how to remove volumes, create snapshots, and set volume bandwidth and IOPS limits.

Adding volumes

Add volumes to a system.

The adding and mapping volume process is necessary, as part of the getting started process, before applications can access the volumes. In addition, you may create additional volumes and map them as part of the maintenance of the virtualization layer.

You can create a volume when the requested capacity is available. To start allocating volumes, the system requires that there be at least three SDS nodes.

Note

For the minimum size of an SDS, see System requirements on page 26.

The created volume cannot be used until it is mapped to (at least) one SDC. For more information, see Mapping a volume to an SDC on page 94.

You can configure the caching option when creating the volumes, or you can change the Read RAM Caching feature later. If you want to enable the caching feature, ensure that the feature is also enabled in the backend of the system, for the corresponding Storage Pool and SDSs. For more information, see Changing RAM Read Cache volume settings on page 197.

Define volume names according to the following rules:

- Contains less than 32 characters
- Contains only alphanumeric and punctuation characters
- Is unique within the object type

ScaleIO objects are assigned a unique ID that can be used to identify the object in CLI commands. You can retrieve the ID via a query, or via the object's property sheet in

the GUI. It is highly recommended to give each volume a meaningful name associated with its operational role.

To add one or multiple volumes, perform these steps:

Procedure

1. In the **Frontend > Volumes** view, navigate to the Storage Pool to which you want to add the volume, and select it.
2. From the **Command menu** or context-sensitive menu, select **Add Volume**.

   The **Add Volume** window is displayed.
3. If you want to create more than one volume, type the number of volumes you would like to add in the **Copies** box.

   a. If you type `1`, only one volume will be created (optional—can be left blank).

   b. If you type a number greater than 1, the characters `%i%` will be added to the **Name** box, and multiple volumes will be created, accordingly.

   The volumes will be named and numbered automatically, starting from 1. If you want the numbering to start from a different number, type it in the **Start numbering at** box, as described in Step 5. The remaining options in the window will be assigned to all the volumes created in this operation.
4. Type a name in the **Name** box.

   If you are adding multiple volumes, they will all be created with the same name, and a number will be appended instead of the characters `%i%`.

   These characters can be positioned anywhere in the name. The names that will be created are displayed in the right pane of the window, as shown in the figure later in this topic.
5. If you want the numbering to start from a specific number other than 1, type it in the **Start numbering at** box.

   This number will be the first number in the series that will be appended to the volume name. For example, if the **Name** is `Vol%i%` and the **Start numbering at** value is `100`, the name of the first volume created will be Vol100, and the second volume will be Vol101, and so on.
6. Type a number in the **Size** box, representing the volume size in GB (basic allocation granularity is 8 GB).
7. Select either **Thick** (default) or **Thin** provisioning options.
8. If obfuscation is required, select the **Use Obfuscation** check box.
9. If you want to disable the RAM Read Cache feature (disabled by default), clear the **Use RAM Read Cache** check box.
10. Click **OK**.

    The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 32 Add Volume window

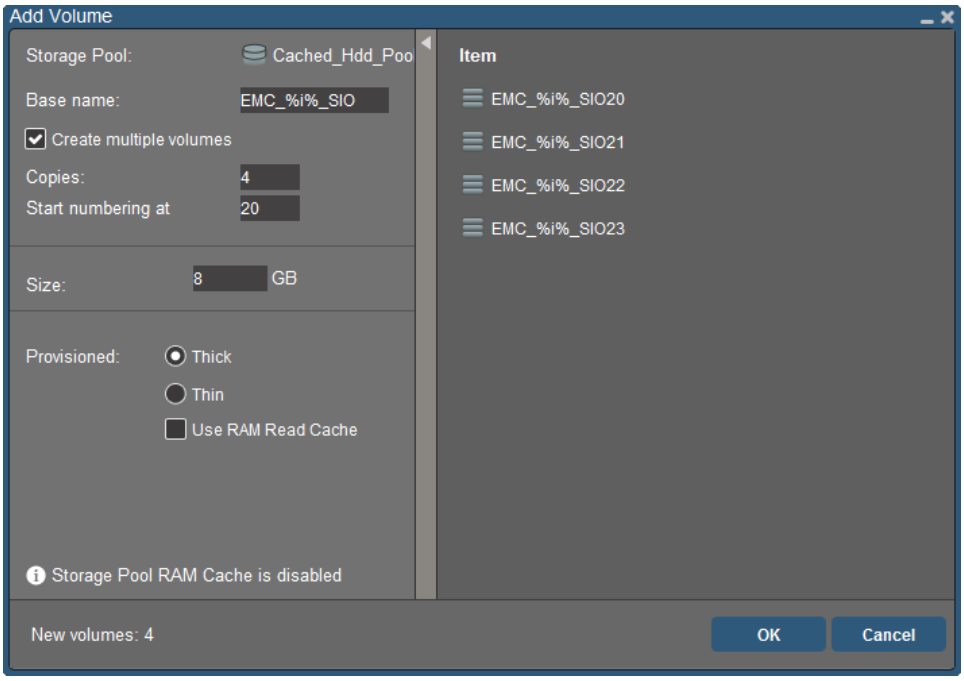

Mapping and unmapping volumes

This topic describes how to map and unmap one or more volumes to/from SDCs. Mapping exposes the volume to the specified SDC, effectively creating a block device on the SDC.

For Linux devices, the `scini` device name can change on reboot. It is recommended to mount a mapped volume to the ScaleIO unique ID, a persistent device name, rather than to the `scini` device name.

To identify the unique ID, run the `ls -l /dev/disk/by-id/` command. For more information, see Associating ScaleIO volumes with physical disks on page 276. You can also identify the unique ID using VMware. In the VMware management interface, device is called **EMC Fibre Channel Disk**, followed by an ID number starting with the prefix **eui**.

To map volumes, perform these steps:

Procedure

1. In the **Frontend** > **Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Map Volumes**.

   The **Map Volumes** window is displayed, showing a list of the volumes that will be mapped.
3. In the **Select Nodes** panel, select one or more SDCs to which you want to map the volumes.
   - You can use the search box to find SDCs.
   - If you select an SDC that is already mapped to the volume, a green icon will appear in the mapping matrix on the right side of the window.
4. Click **Map Volumes**.

The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 33 Map Volumes window after mapping is complete

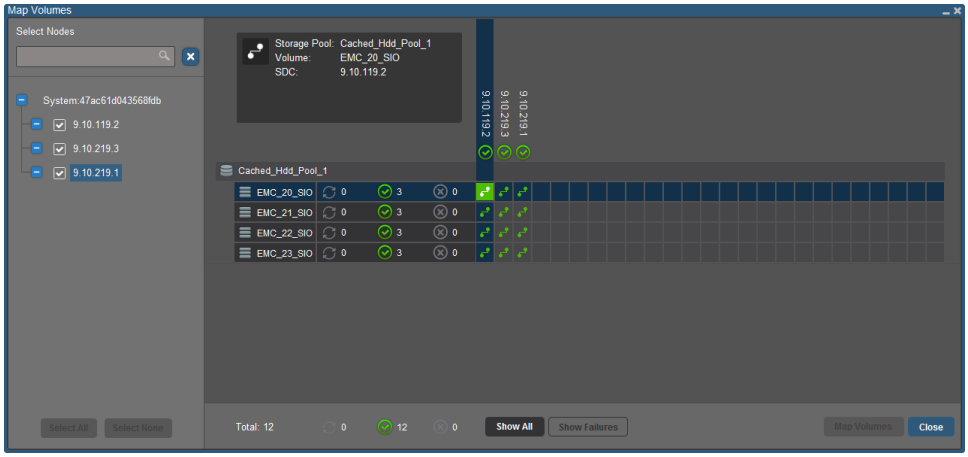

To unmap volumes, perform these steps:

5. In the **Frontend > Volumes** view, navigate to the volumes, and select them.
6. From the **Command menu** or context-sensitive menu, select **Unmap Volumes**.

   The **Unmap Volumes** window is displayed, showing a list of the volumes that will be unmapped.
7. If you want to exclude some SDCs from the unmap operation, in the **Select Nodes** panel, select one or more SDCs for which you want to retain mapping.

   You can use the search box to find SDCs
8. Click **Unmap Volumes**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 34 Unmap Volumes window

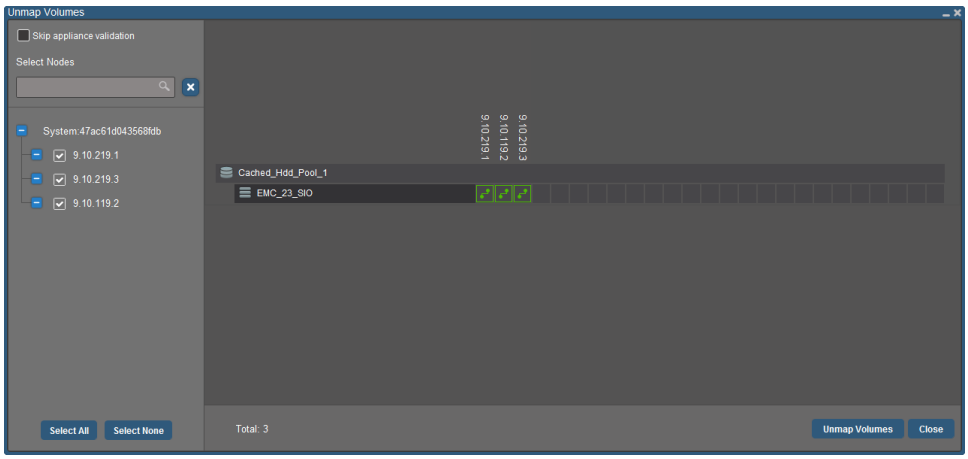

Removing volumes

Before removing a volume from a system, you must ensure that it is not mapped to any SDCs. If it is, unmap it before removing it. For information, see Mapping and unmapping volumes on page 202.

If you want to remove a volume's related snapshots as well, or just the snapshots, see Removing snapshots on page 206. Before removing snapshots, you must unmap all of them before removing them, in the same way that you unmap volumes.

Best practice is to avoid deleting volumes or snapshots while the MDM cluster is being upgraded, to avoid causing a Data Unavailability status.

Note

Removal of a volume erases all the data on the corresponding volume.

To remove one or multiple volumes, perform these steps:

Procedure

1. In the **Frontend** > **Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Remove**.

   The **Remove Volumes** window is displayed, showing a list of the volumes that will be removed.
3. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

   Figure 35 Remove Volumes window

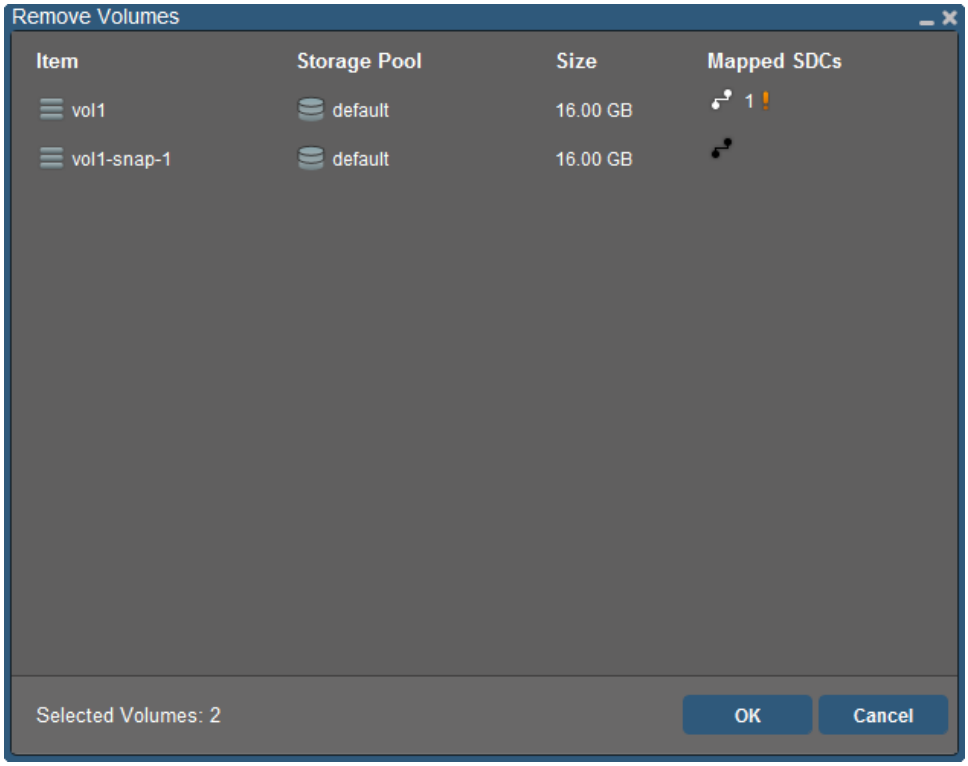

Creating volume snapshots

This topic describes how to take a snapshot of one or more volumes.

When specifying more than one volume (a list), a consistency group is generated by default, and can be viewed in the snapshot's property sheet. The snapshots under the consistency group are taken simultaneously for all listed volumes, thus ensuring their consistency.

You can accept the default name, or define snapshot names according to the following rules:

- Contains less than 32 characters
- Contains only alphanumeric and punctuation characters
- Is unique within the object type

ScaleIO objects are assigned a unique ID that can be used to identify the object in CLI commands. You can retrieve the ID via a query, or via the object's property sheet in the GUI.

Note

The consistency group is for convenience purposes ONLY. There are no protection measures to conserve the consistency group. You can delete members from it.

To take a snapshot, perform these steps:

Procedure

1. In the **Frontend** > **Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Snapshot Volume**.

   The **Snapshot Volume** window is displayed, showing the volumes for which snapshots will be created.
3. In the **Index** box, type the number that you want to append to the snapshot names.
4. If you want the snapshots to belong to a consistency group, ensure that the **Create Consistency Group** check box is selected.
5. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 36 Snapshot Volume window

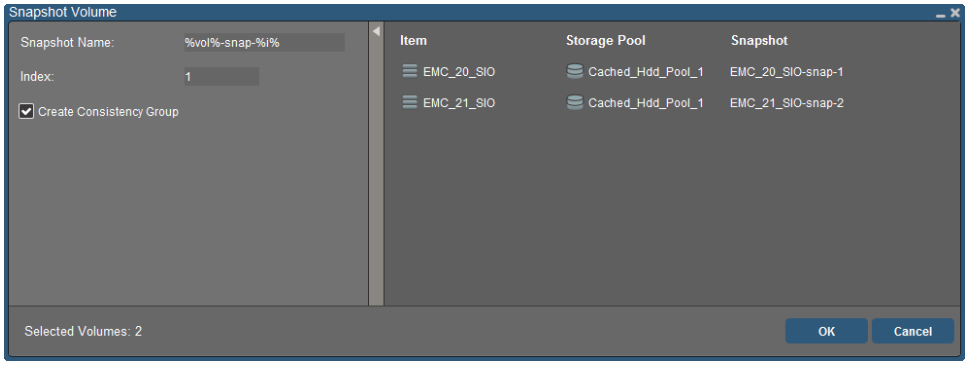

Removing snapshots

This topic explains how to remove a volume together with its snapshots, or remove snapshots only. Before removing a volume or snapshots, you must ensure that they are not mapped to any SDCs. If they are, unmap them before removing them. Snapshots are unmapped in the same way as volumes are unmapped. For information, see Mapping and unmapping volumes on page 202.

Best practice is to avoid deleting volumes or snapshots while the MDM cluster is being upgraded, to avoid causing a Data Unavailability status.

Note

Removal of a volume or snapshot erases all the data on the corresponding volume or snapshot.

To remove snapshots, perform these steps:

Procedure

1. In the **Frontend** > **Snapshots** view, navigate to the volume from which you want to remove snapshots, and select it.
2. From the **Command menu** or context-sensitive menu, select one of the following options, depending on your needs:
   a. To retain the parent volume, and remove only its snapshots, select **Remove Descendants Only**
   b. To remove the both the parent volume and all volumes that were created as snapshots of the specified volume or one of its descendants, select **Remove with Descendants**

      The **Remove Volumes** window is displayed, showing a list of the objects that will be removed.
3. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 37 Removal of a volume and associated snapshots

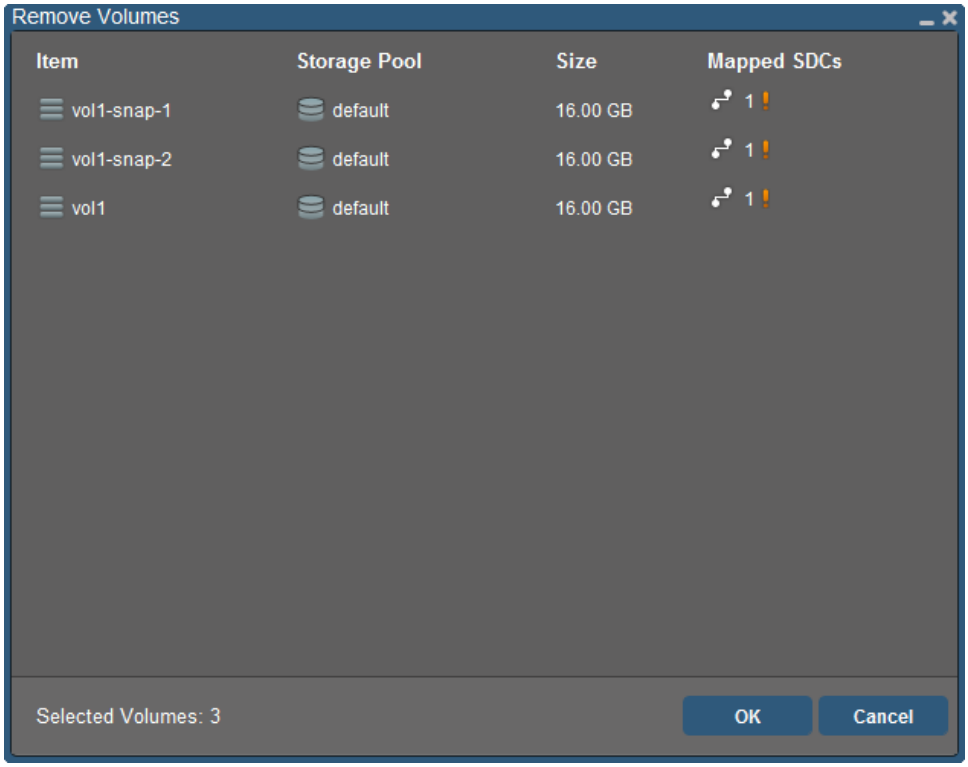

Figure 38 Removal of snapshots of a specific volume

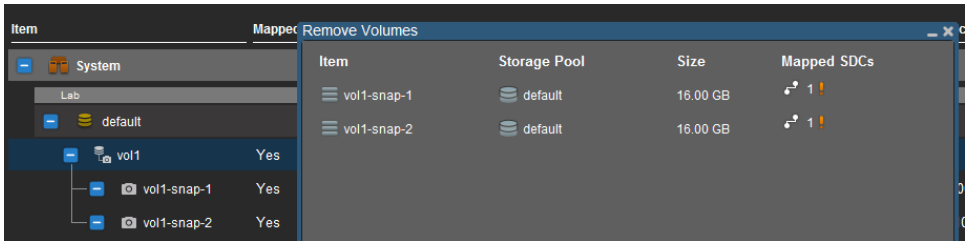

Removing snapshots from a consistency group

To remove a snapshot from a consistency group, perform these steps:

Procedure

1. In the **Frontend** > **Snapshots** view, navigate to the snapshot which you want to remove from the consistency group, and select the snapshot.
2. From the **Command menu** or context-sensitive menu, select **Remove Consistency Group**.

   The **Remove Consistency Group** window is displayed, showing the selected snapshot.
3. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 39 Remove Consistency Groups window

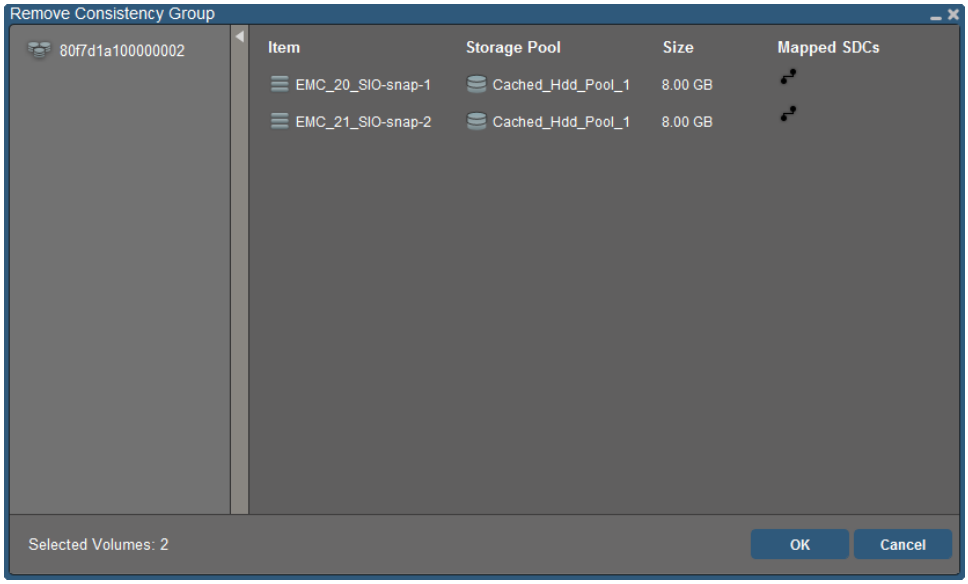

Increasing a volume's size

This topic describes how to increase the size of one or more volumes in the system. You can increase (but not decrease) a volume capacity at any time, as long as there is enough capacity for the volume size to grow.

To increase the size of the specified volumes, perform these steps:

Procedure

1. In the **Frontend > Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Increase Volumes' Size**.

   The **Increase Volumes' Size** window is displayed, showing a list of the volumes that will be modified.
3. In the **New Size** box, type a number representing the new volume size in GB (basic allocation granularity is 8 GB).
4. Click **OK**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 40 Increase Volumes' Size window

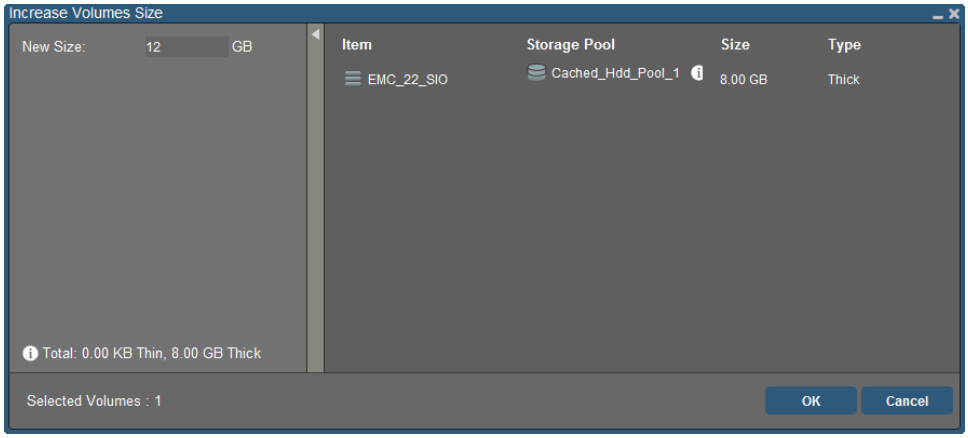

Setting volume bandwidth and IOPS limits

This topic describes how to set bandwidth and IOPS limits for volumes. The limits will be applied on a per SDC basis. This enables you to control the quality of service (QoS). Ensure that the volumes are mapped before you set these limits.

To set limits on volumes, perform these steps:

Procedure

1. In the **Frontend > Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Set Volume Limits**.

   The **Set Volume Limits** window is displayed, showing a list of the volumes that will be modified.
3. In the **Bandwidth Limits** and **IOPS Limits** boxes, type the required values, or select the corresponding **Unlimited** check box.
   - The number of IOPS must be larger than 10.
   - The volume network bandwidth is in MB/sec.
4. In the **Select Nodes** panel, select the SDCs to which you want to apply the changes.
5. Click **Set Limits**.

   The progress of the operation is displayed at the bottom of the window. It is recommended to keep the window open until the operation is completed, and until you can see the result of the operation.

Figure 41 Set Volume Limits window

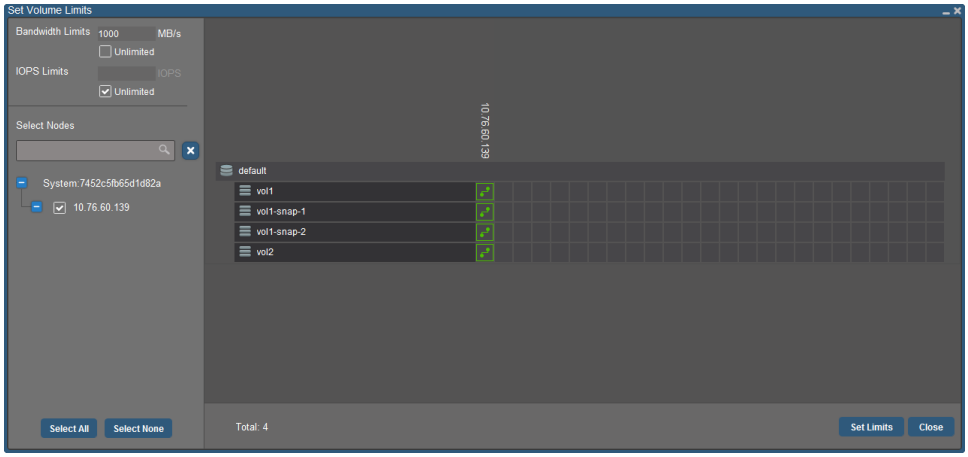

Entering and exiting SDS Instant Maintenance Mode

This topic explains how to put an SDS into Maintenance Mode, in order to perform non-disruptive maintenance on the SDS, and how to cancel Maintenance Mode when you are finished. Instant Maintenance Mode lets you restart a server that hosts an SDS, without initiating data migration or exposing the system to the danger of having only a single copy of data. The system displays the SDSs that are in Maintenance Mode at any given time (but does not provide the total number of SDSs).

While SDSs are in Maintenance Mode, you should avoid both unnecessary rebuilds and operations that require taking SDS offline temporarily. It is recommended to use Maintenance Mode when there is relatively low system activity, as the time it takes for an SDS to exit Maintenance Mode depends on the amount of data that needs to be synchronized back into the server.

Note

Functional operations, such as configuration, cannot be performed on an SDS while it is in Maintenance Mode.

If an active full copy is lost, its data will be unavailable until the SDS is brought back into the system, but that data will not be lost; it will be rebuilt using the temporary copy.

To place an SDS in Maintenance Mode, perform these steps:

Procedure

1. In the **Backend** view, navigate to, and select the desired SDS.
2. From the **Command menu** or context-sensitive menu, select **Enter Maintenance Mode**.

   The **Enter Maintenance Mode** window is displayed.
3. If you want to force entry into Maintenance Mode even though there is insufficient space or degraded/failed capacity, select the corresponding check box:
   - **Force Insufficient Space**—allow entry into maintenance mode, even without enough available capacity
   - **Force Degraded or Failed**—allow entry into maintenance mode, even with degraded or failed data

4. Click **OK**.

   The status area at the bottom of the window indicates when the operation is complete. Once the SDS is in Maintenance Mode, this will be indicated both on the Dashboard, and in Backend tables and Property Sheets, using the symbol, and the Maintenance Mode color code (green). To put an SDS back into regular service (cancel Maintenance Mode), perform these steps:

5. In the **Backend** view, navigate to, and select the desired SDS.
6. From the **Command menu** or context-sensitive menu, select **Exit Maintenance Mode**.

   The **Exit Maintenance Mode** window is displayed.

7. If you want to force exit from Maintenance Mode even though there is a failed SDS, select the **Force Failed SDS** check box.
8. Click **OK**.

   The status area at the bottom of the window indicates when the operation is complete. Once the operation has been successfully completed, the SDS returns to normal operation, and data deltas collected on other SDSs during the maintenance period are copied back to the SDS.

Configuring and viewing Oscillating Failure counters

Oscillating failure handling provides the ability to handle error situations, and to reduce their impact on normal system operation. This feature detects and reports various oscillating failures, in cases when components fail repeatedly and cause unnecessary failovers. You can configure the time interval associated with each window type, and the number of failures allowed before reporting commences for each window type, per counter.

You can reset specified oscillating failure counters to zero. This can be useful when you have fixed a problem and want to ensure that an alert is no longer active in the system.

Configuring Oscillating Failure counter parameters

Configure Oscillating Failure counter parameters for the entire system, or for specific Protection Domains or Storage Pools.

Procedure

1. Perform one of the following:

Option	Description
To configure counter parameters for all SDCs, Protection Domains or Storage Pools in the system:	In the **Backend > Storage** view, select the **System** icon.
To configure counter parameters for a specific Protection Domain or Storage Pool:	In the **Backend > Storage** view, navigate to, and select the desired Protection Domain or Storage Pool.

2. From the **Command menu** or context-sensitive menu, select **Set Oscillating Failure Properties**.

3. Perform one of the following:

Option	Description
For system level:	In the **For All** box, select an option: **SDCs**, **Protection Domains**, or**Storage Pools**.
For a Protection Domain or a Storage Pool:	Go to the next step.

4. In the **Counter Type** box, select a counter. Options vary, depending on the item selected in the previous step.
5. In the **Window Type** box, select an option for the sliding window interval: **Short**, **Medium** or **Long**.
6. Perform one of the following:

Option	Description
If you want to remove the selected counter definition from the system:	Select the **Remove the counter** check box.
If you want to modify the threshold for the selected counter definition:	Enter a number in the fields for: • **failures** (the maximum number of failures per time interval before reporting begins) • **seconds** (the number of seconds per time interval)

7. Click **OK**.

   The currently configured counter parameters are displayed in the corresponding Property Sheet, in the **Oscillating Failure Parameters** section.

Figure 42 Configure Oscillating Failure counters—System

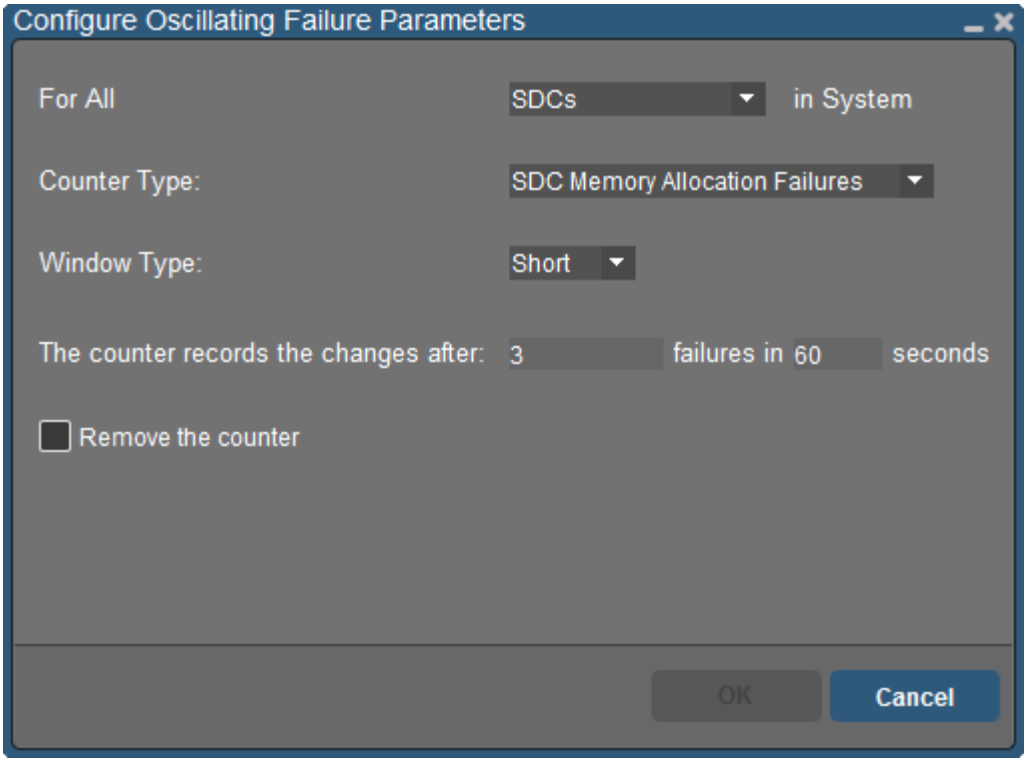

Figure 43 Configure Oscillating Failure counters—Protection Domain or Storage Pool

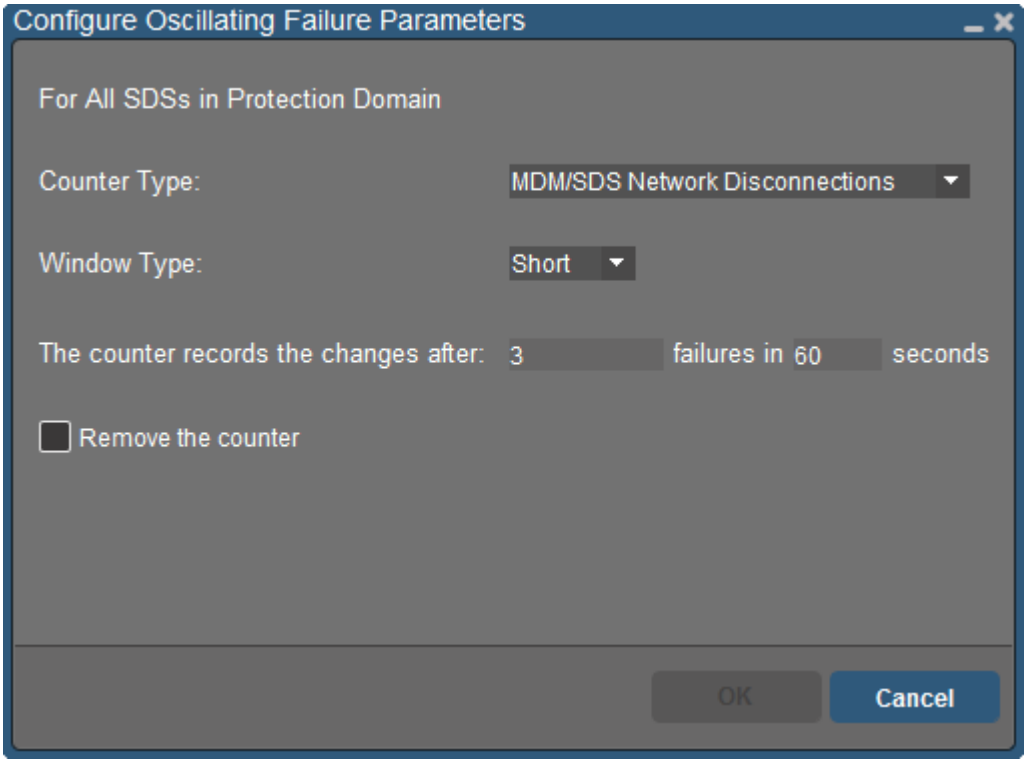

Resetting Oscillating Failure counters

You can reset specified oscillating failure counters to zero. This can be useful when you have fixed a problem and want to ensure that an alert is no longer active in the system. You can reset counters for the entire system, per Protection Domain, or per Storage Pool.

To reset oscillating failure counters for all SDCs, Protection Domains or Storage Pools in the system, perform the following steps:

Procedure

1. In the **Backend** view, select the System icon.
2. From the **Command menu** or context-sensitive menu, select **Reset Oscillating Failure Counters**.
3. In the **For All** box, select an option.

   If you want to reset counters for all object types, select **Objects**.
4. Do one of the following:
   a. For a specific counter, in the **Counter Type** box, select the required counter.
   b. For all counters, in the **Counter Type** box, select **None**, and select the **Reset All Counters** check box.
5. Click **OK**.

   Figure 44 Reset Oscillating Failure counters—System

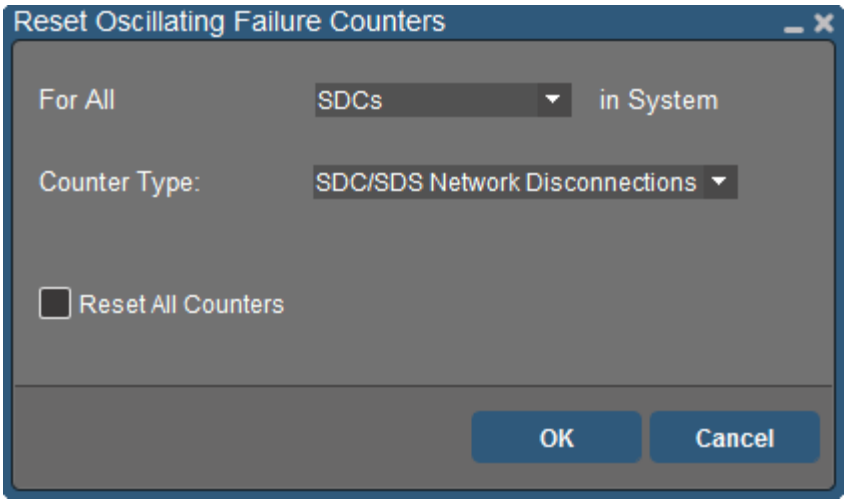

   To reset counters for a specific Protection Domain or Storage Pool, perform the following steps:

6. In the **Backend** view, navigate to, and select the desired Protection Domain or Storage Pool.
7. From the **Command menu** or context-sensitive menu, select **Reset Oscillating Failure Counters**.
8. Do one of the following:
   a. For a specific counter, in the **Counter Type** box, select the required counter.

b. For all counters, select the **Reset All Counters** check box.

9. Click **OK**.

Figure 45 Configure Oscillating Failure counters—Protection Domain or Storage Pool

Viewing Oscillating Failure counters

You can view Oscillating Failure counters for network related issues, for SDCs, for SDSs and for devices in the following ways:

- Network:

1. In the **Backend** view, select the **System** icon.
2. From the **Command menu** or context-sensitive menu, select **Download Network Failure Counters**.
3. Browse to the location in which the file will be saved, and click **OK**. A JSON file containing the counters is saved in the location that you specified.
   - In Windows, view the file in a text editor, such as Notepad++.
   - In Linux, use the `more` command to view the file (for example, `more Oscillating_Network_Failures_Counters_2016-01-28-13-31-57.json`)

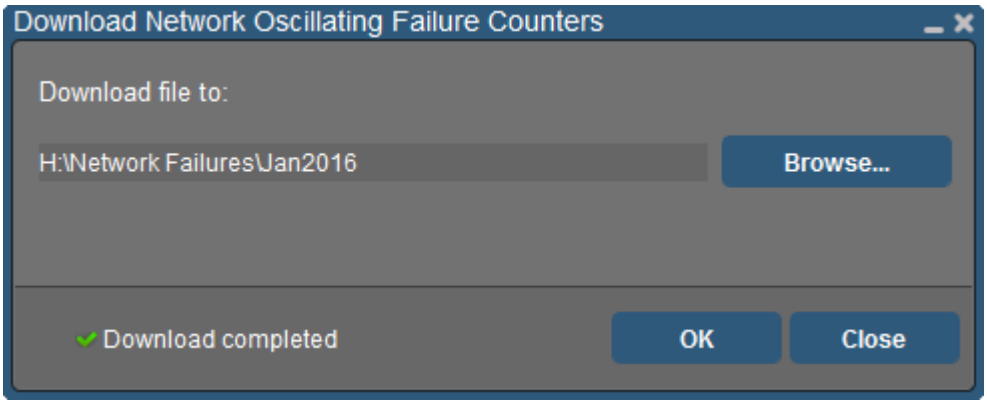

- SDCs:

1. In the **Frontend** view, select the **SDCs** option, and then select the required SDC.
2. Open the Property Sheet, and click the **Oscillating Failure Counters** section.

Note

When there are no oscillating failures counters for SDC, the **Oscillating Failures Counters** section displays `None Found`.

- SDSs and devices:

Procedure

1. In the **Backend** view, navigate to and select the required SDS or device.
2. Open the Property Sheet, and click the **Oscillating Failure Counters** section.
   a. If there is an active alert, you can view the oscillating failures alert in the **Alerts** section of the Property Sheet.

   Note

   When there are no oscillating failures counters for SDS or device, the **Oscillating Failures Counters** section displays `None Found`.

Applying Performance Profiles to system components

You can use the GUI to apply performance profiles to system components. The high performance profile configures a predefined set of parameters for very high performance use cases. When a container is provided in the command (System/Protection Domain/Fault Set), all the objects currently in the container are configured.

Note

For a complete list of parameters controlled by the profiles, refer to the document *ScaleIO v2.0.x Performance Fine-Tuning Technical Notes*.

The profiles are applied separately to:

- SDSs
- SDCs
- MDM cluster

Note

After changing the performance profile of an SDS (on an SVM), you must perform manual memory allocation on the SVM, as described in the *ScaleIO Deployment Guide*.

To apply a profile to system components, perform the following steps:

Procedure

1. In the **Backend** or **Frontend** view, navigate to, and select the desired objects.

   Note

   If you want to apply the Performance Profile to MDMs, select the System object.

2. From the **Command menu** or context-sensitive menu, select **Set Performance Profile for XXX**, where XXX represents one of the following:
   - MDMs
   - All SDSs
   - SDS
   - All SDCs
   - SDC

   The **Set Performance Profile for** window is displayed.
3. Select a profile from the drop-down list, and click **OK**.

   Figure 46 Set Performance Profile window

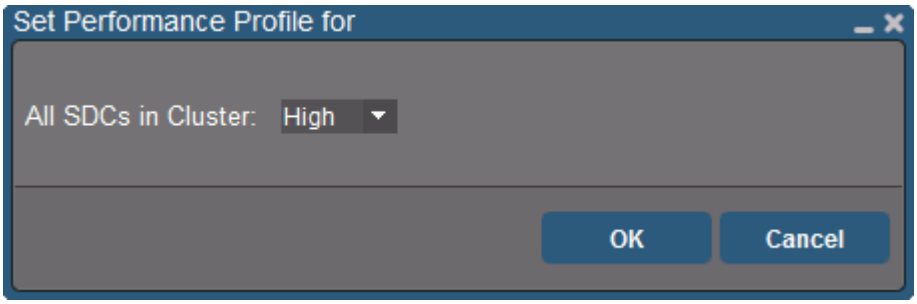

Configuring I/O priorities and bandwidth use (advanced)

The ScaleIO system includes advanced settings which control I/O priorities and bandwidth use, which can be used to fine-tune system performance. It is recommended to retain default settings, unless you are an advanced user.

Application IOPS and bandwidth (advanced)

Priority can be given to different types of I/Os in the system, including application I/Os. The number of concurrent Rebuild and Rebalance jobs can be configured, along with bandwidth used for these jobs. I/O prioritization is configured per Storage Pool.

NOTICE

These features affect system performance, and should only be configured by an advanced user.

Give priority to Application I/Os during Rebuild and Rebalance jobs, by performing these steps:

Procedure

1. In the **Backend** view, navigate to, and select the desired Storage Pool.
2. From the **Command menu** or context-sensitive menu, select **Set I/O Priority**.
3. Select Favor Application I/O for Rebalance and Rebuild, and click **OK**.

System IOPS and bandwidth (advanced)

You can prioritize different types of I/O in the system, as well as configure network throttling per Protection Domain.

I/O prioritization

Priority can be given to different types of I/Os in the system. The number of concurrent Rebuild and Rebalance jobs can be configured, and bandwidth for Rebalance jobs can be configured. If the `Dynamic Bandwidth Throttling` option is selected, additional items can be configured, such as `Application IOPS threshold`, `Application bandwidth threshold`, and `Application threshold quiet period`. Default values for these features are provided in the *ScaleIO CLI Reference Guide*.

NOTICE

These features affect system performance, and should only be configured by an advanced user.

Configure I/O prioritization for Rebuild and Rebalance by performing these steps:

Procedure

1. In the **Backend** view, navigate to, and select the desired Storage Pool.
2. From the **Command menu** or context-sensitive menu, select **Set I/O Priority**.
3. Select the desired options and edit values, and click **OK**.

Configuring Network Throttling

Network throttling affects network limits, and is used to control the flow of traffic over the network. It is configured per Protection Domain. The SDS nodes transfer data between themselves. This data consists of user-data being replicated as part of the RAID protection, and data copied for internal rebalancing and recovery from failures. You can modify the balance between these types of data loads by limiting the data copy bandwidth. This change affects all SDSs in the specified Protection Domain.

NOTICE

These features affect system performance, and should only be configured by an advanced user. Contact EMC support before you change this configuration.

Procedure

1. In the **Backend** view, navigate to the desired Protection Domain, and select its row in the table.
2. From the **Command menu** or context-sensitive menu, select **Set Network Throttling**.

   The **Set Network Throttling** window is displayed.
3. Configure the settings, and click **OK**.

Figure 47 Set Network Throttling window

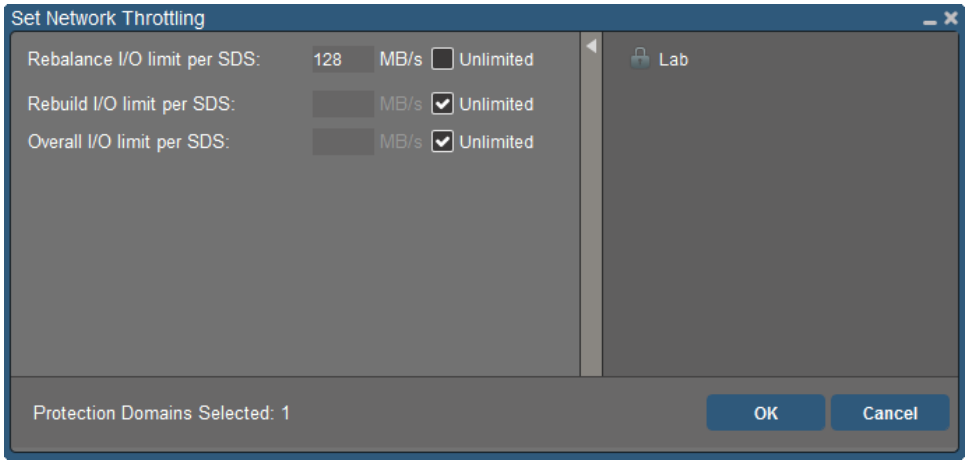

Enabling and disabling Rebuild\Rebalance (advanced)

By default, Rebuild and Rebalance features are enabled in the system, because they are essential for system health, optimal performance, and data protection. These features should only be disabled temporarily in very specific circumstances, and should not be left disabled for long periods of time. Rebuild and Rebalance features are enabled and disabled per Storage Pool.

NOTICE

Rebuilding is an essential part of the ScaleIO system, which provides protection for your data. It is not recommended to disable the Rebuild feature, except in very special circumstances. Rebalancing is an essential part of the ScaleIO system, and should only be disabled, temporarily, in special circumstances. Disabling rebalance may cause the system to become unbalanced even if no capacity is added or removed. For example, during a recovery from an SDS or device failure, some rebalance activity may be needed to ensure optimal balancing.

To enable or disable Rebuild and Rebalance features, perform these steps:

Procedure

1. In the **Backend** view, navigate to, and select the desired Storage Pools.
2. From the **Command menu** or context-sensitive menu, select **Enable/Disable Rebuild/Rebalance**.

   The **Enable or Disable Rebuild and Rebalance** window is displayed.
3. Select or clear the options that you require (selected=enable; clear=disable), and click **OK**.

Using the background device scanner

The background device scanner scans devices in the system to check for errors. You can enable and disable the background device scanner, as well as reset the background device scanner counters. Information about errors will be provided in event reports. For more information about viewing events, see Viewing events on page 295.

Enabling and disabling the background device scanner

The scanner can be enabled on all the devices in the specified Storage Pool. There are two modes: **device only** mode, and **data comparison** mode:

- Device only—Perform read operations. Fix from peer on errors.
- Data comparison—Perform the device-only test, and compare the data content with peer. Zero padding must be enabled in order to set the background device scanner to data comparison mode.

To enable or disable the background device scanner, follow these steps:

Procedure

1. In the **Backend** view, navigate to, and select the desired Storage Pools.
2. From the **Command menu** or context-sensitive menu, select **Set Background Device Scanner Mode**.

   The **Configure Storage Pool Background Device Scanner** window is displayed. The right pane of the window displays the Storage Pools that you are configuring.
3. For the **Enable Background Device Scanner** option, do one of the following:
   a. To enable the scanner, select the check box, and proceed to the next step.
   b. To disable the scanner, clear the check box, and click **OK** to finish.
4. Select an option (selected=enable; clear=disable):
   a. **Device only**
   b. **Data comparison**
5. In the **Bandwidth Limit** box, accept the default or type a number in KB per second (per device).

   The given value should be in the range 10 KB-10 MB (default = 1 MB).

   Note

   High bandwidth should be used very carefully for extreme cases only (such as an urgent need to check some devices), as it may create negative impact on system performance. Setting the background device scanner bandwidth should take into account maximum bandwidth of the devices.

6. Click **OK**.

Figure 48 Background device scanner configuration

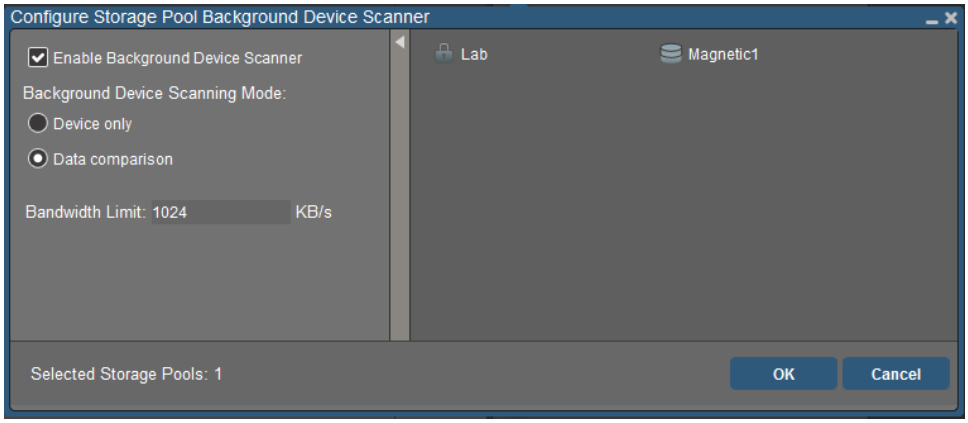

Resetting the background device scanner counters

You can reset background device scanner error counters for specified Storage Pools. Counters for data comparison errors, or corrected read errors, or both counter types can be reset.

To reset counters, follow these steps:

Procedure

1. In the **Backend** view, navigate to, and select the desired Storage Pools.
2. From the **Command menu** or context-sensitive menu, select **Reset Background Device Scanner Counters**.

   The **Reset Background Device Scanner Counters** window is displayed. The right pane of the window shows the Storage Pools that you are configuring.
3. Select or clear the option that you require, or both options (selected=enable; clear=disable).
4. Click **OK**.

   Figure 49 Reset Background Device Scanner Counters window

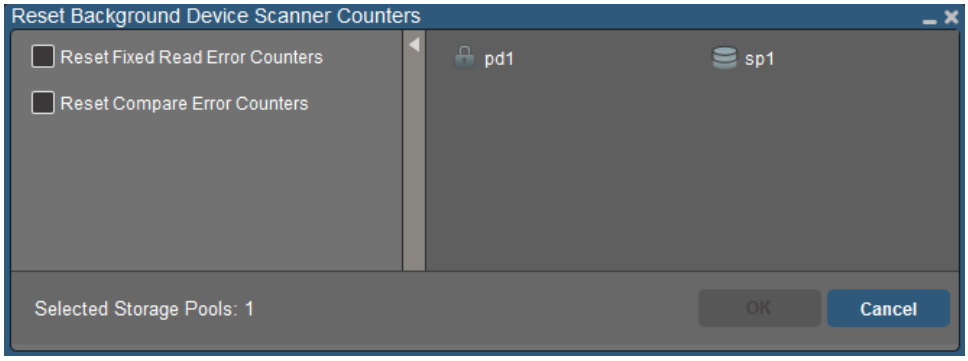

Modifying Checksum protection mode

Checksum mode can be used to validate in-flight data reads and writes, in order to protect data from data corruption. To modify this setting, perform the following steps:

Procedure

1. In the **Backend** view, navigate to, and select the desired Storage Pools.
2. From the **Command menu** or context-sensitive menu, select **Configure Use Checksum**.

   The **Configure Use Checksum** window is displayed.
3. Do one of the following:
   a. To enable the Checksum feature, select the **Enable Use Checksum** check box.
   b. To disable the Checksum feature, clear the **Enable Use Checksum** check box.
4. Click **OK**.

   Figure 50 Configure Use Checksum window

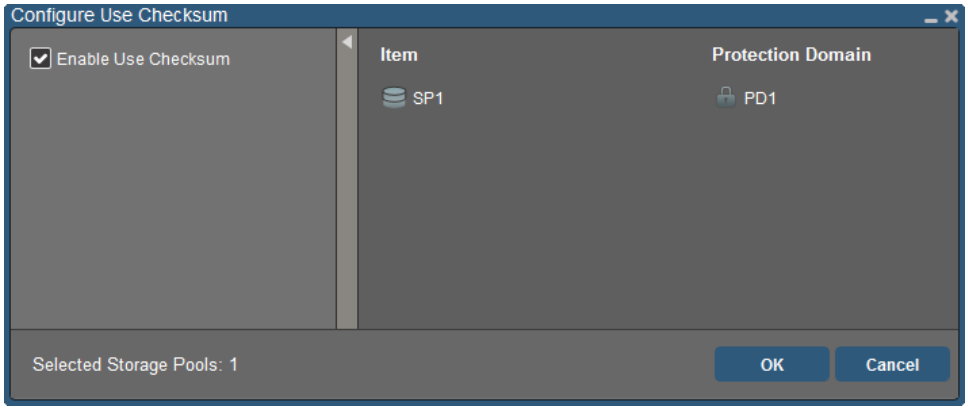

Renaming objects

Object names are used to identify the objects in the GUI, and can also be used to specify objects in CLI commands. You can view an object's name in its Property Sheet, in the **Identity** section.

Note

It is not possible to rename a Read Flash Cache device using this command.

You can define object names according to the following rules:

1. Contain less than 32 characters
2. Contain only alphanumeric and punctuation characters
3. Be unique within the object type

When a name has not been defined, the system may display default system-defined names, as follows:

- SDC—its first IP address
- SDS—its first IP address
- Device—the path to the device
- All other objects—the object's ID

Note

A name must be assigned to a volume when it is initially created. You can rename the volume later, using the **Rename** command.

Procedure

1. Depending on the object type, in the **Backend** or **Frontend** view, navigate to the object in the table, and select its row.
2. From the **Command menu** or context-sensitive menu, select the **Rename** command.

   An editing window is displayed, showing the current name, and an editable field for the new name.
3. Type the new name in the field, and click **OK**.

Approving pending security certificates

You can approve pending security certificates, and view approved certificates in the **System Settings** window, by following these steps:

Procedure

1. From the **Preferences** menu in the top right corner, select **System Settings**.

   The **System Settings** window is displayed.
2. In the left pane, click **Certificates**.

   Approved and pending certificates are displayed.
3. Scroll to the **Pending approval** section, and expand the section.
4. Scroll to the bottom of the information about the required certificate, and click **Confirm**.

   Note

   When there are pending certificates in the system, they are listed in the **Backend** view > **State Summary** table, and in the **Alerts** view.

Customizing system preferences

You can customize various features in the GUI using the **User Preferences** window. The following features can be customized:

- Refresh data rate
- Clear host history from previous sessions
- Calculation of I/O workload average rate shown on the Dashboard
- System clock display
- Advanced display mode for Dashboard, Backend internal views, and Property Sheet
- Log level

Procedure

1. From any location in the GUI, open the **Preferences** menu in the top right corner, and select **User Preferences**.

   The **User Preferences** window is displayed.

2. Edit the options according to your needs, and click **Apply**.

   Figure 51 Preferences window

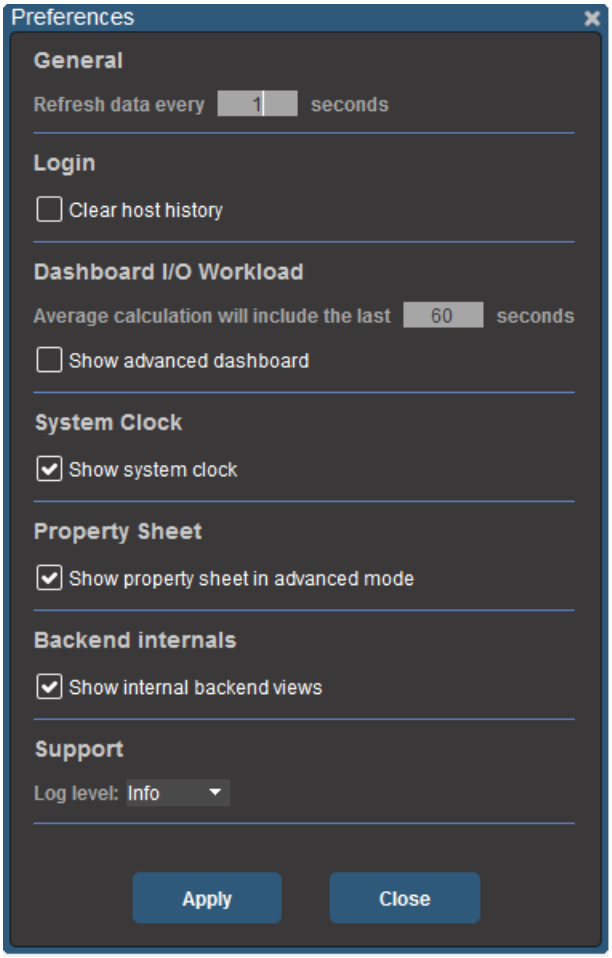

Table 29 Preferences

Item	Description
Refresh data every *n* seconds	Controls the rate at which data displayed in the GUI is refreshed, in seconds (Default: 10 seconds) The refresh occurs at least at the specified rate. It is not intended to be used as a means of limiting client traffic, although it would actually do so.
Clear host history	When selected, the GUI does not save and present host connection details from previous sessions
Dashboard I/O workload: average calculation will include the last *n* seconds	Controls the time period used when averages are computed and displayed by the GUI (default: 10 seconds)

Table 29 Preferences (continued)

Item	Description
Show advanced dashboard	When selected, (default), includes more details on some tiles in the Dashboard view. The toggle buttons switch between the statistics displayed in large fonts and small fonts. The upper button toggles between average values and sample values. The lower button toggles between display of bandwidth or IOPs in large fonts on this tile. The symbol means that the number displayed is the average taken during the last *n* seconds. *n* can be configured in **Dashboard I/O Workload** in this window. The symbol means that the number displayed is from the last data sample that was taken. The period between automatic refreshes can be configured in **Dashboard I/O Workload** in this window.
Show Property Sheet in advanced mode	Displays additional details in Property Sheets: • Capacity section—Snapshot Capacity Reserved • Rebuild/Rebalance—Data Movement Jobs • RAM Read Cache—Cache Evictions, Cache Entry, and Cache Skip tables These details are usually only relevant for advanced users and technical support purposes.
Show internal Backend views	Displays additional options for Backend table views. These options are recommended only for advanced users and technical support purposes.
Log level	Controls the type of data saved in system logs, which may be required by Customer Support for troubleshooting purposes. The default setting recommended for regular operation is Info. Other options include: Trace, Debug, Warn, and Error. **Note** Trace and Debug options may affect system performance, and are usually only recommended for technical support purposes. For more information about logs, see your system's Log Collection Technical Notes.

CHAPTER 13

Using the VMware Plug-in

This chapter describes how to use the VMware plug-in (the “plug-in”) to view and provision ScaleIO components. Topics include:

- VMware Plug-in overview 228
- Configuring components 229
- Viewing components 238

VMware Plug-in overview

The VMware plug-in communicates with the MDM and the vSphere server, enabling you to view components and perform many configuration/provisioning tasks right from within the VMware environment.

Before benefiting from ScaleIO, you must create volumes and map them to SDCs installed on the ESX hosts. This requires the following steps:

1. Creating a volume
2. Mapping the volume to all SDCs

This set of manual tasks is automated in the plug-in, as described in "Creating, mapping, and unmapping volumes".

To use the plug-in, it must be registered in your vCenter. For more information, see the *ScaleIO Deployment Guide*.

To open the plug-in, in the vSphere Web Client click .

The EMC ScaleIO screen appears.

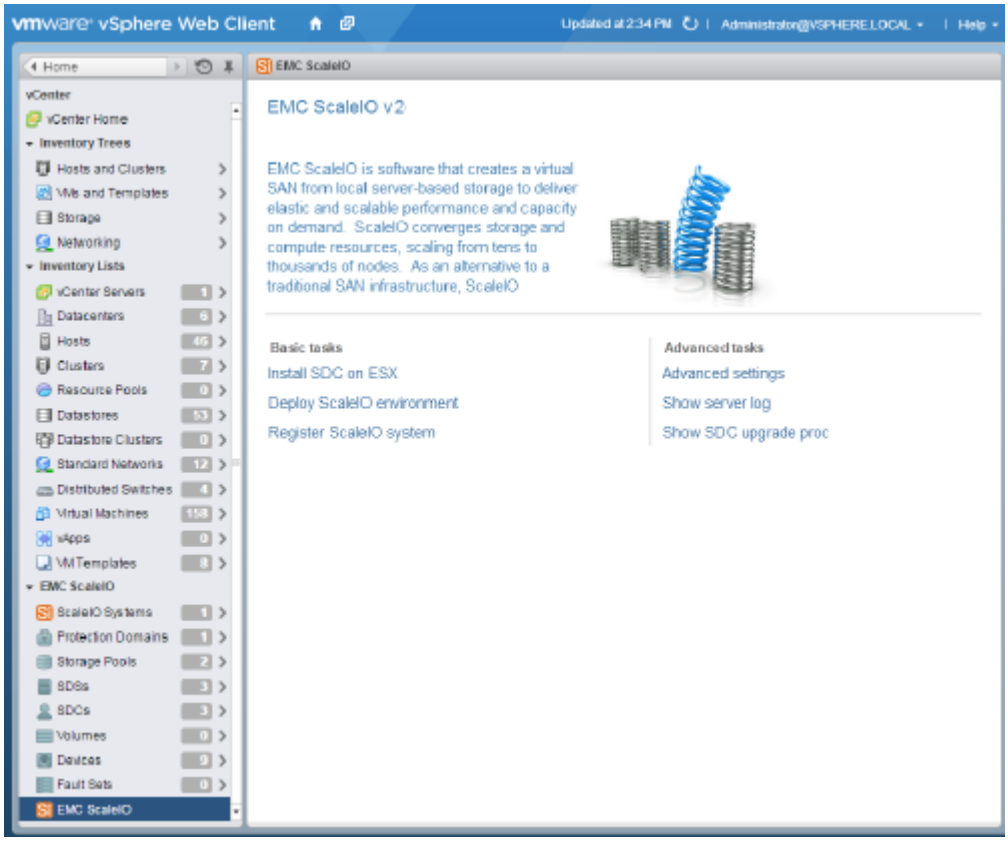

The ScaleIO screen displays an overview of the configured components. In this example, one system has been configured, with the following components:

- Protection Domain—1
- Storage Pool—2
- SDS—3
- SDC—3

- Volumes—0
- Devices—9
- Fault Set—0

You can use the plug-in to configure and view ScaleIO components.

Configuring components

There are two levels of component configurations:

- Basic: The basic configurations are all performed the same way. The process is described just once.
- Advanced: Each advanced configuration setting has a unique dialog box, which is described in "Configuring components—advanced".

The following table lists the activities you can perform and categorizes each as basic or advanced:

Table 30 Plug-in activity matrix

Object	Perform this activity	Basic or advanced	Access from this screen
System	Deploy ScaleIO system	Advanced. See the EMC ScaleIO Deployment Guide.	EMC ScaleIO
	Register an existing system	Basic. Enter the system Master MDM IP address, user name, and password.	EMC ScaleIO
	Unregister a system	Basic	ScaleIO Systems
	Update system credentials	Basic. Enter new user name and password.	ScaleIO Systems
	Configure virtual IPs	Advanced	ScaleIO Systems
ScaleIO Gateway	Register/ Update Gateway	Basic. Enter IP address, OS user name, and OS password.	ScaleIO Systems
	Open Gateway	Basic. Navigates to the Gateway Installation Manager.	ScaleIO Systems
Protection Domain	Create a Protection Domain	Basic	ScaleIO Systems
	Remove a Protection Domain	Basic	Protection Domains
Storage Pool	Create a Storage Pool	Basic	Protection Domains
	Remove a Storage Pool	Basic	Storage Pools

Table 30 Plug-in activity matrix (continued)

Object	Perform this activity	Basic or advanced	Access from this screen
	Configure RAM Read Cache	Basic[a]	Storage Pools
SDS	Add a device to an SDS	Advanced	SDSs
	Remove a device from an SDS	Basic	Devices
SDC	Install SDC on ESX	Advanced. See the *ScaleIO Deployment Guide*.	EMC ScaleIO
	Upgrade SDC	Advanced	SDCs
	Update SDC Parameters	Advanced. See the *ScaleIO Deployment Guide*.	EMC ScaleIO
Volume	Create and map volumes	Advanced	Storage Pools
	Map a volume	Advanced	Volumes
	Remove a volume (must be unmapped first)	Basic	Volumes
	Unmap a volume	Advanced	Volumes
	Configure RAM Read Cache	Basic[a]	Volumes
Fault Set	Create a Fault Set	Basic[b]	Protection Domains
Device	Clear a device error **Note** Removes the error message. Can be performed only after clearing the error.	Basic	Devices
	Add a device to an SDS	Advanced	SDSs
	Remove a device from an SDS	Basic	Devices

Table 30 Plug-in activity matrix (continued)

a. For RAM Read Cache to work on a volume, both the volume and its Storage Pool must have the feature enabled.
b. When defining Fault Sets, you must follow the guidelines described in Fault Sets on page 38. Failure to do so may prevent creation of volumes.

Configuring components—basic

This section shows how to perform basic configuration activities. All activities are performed from the **Actions** menu in each screen and entering simple information.

For example, to create a Protection Domain from the **ScaleIO Systems** screen, perform the following:

Procedure

1. From the **ScaleIO Systems** screen, click **Actions** > **Create Protection Domain**:

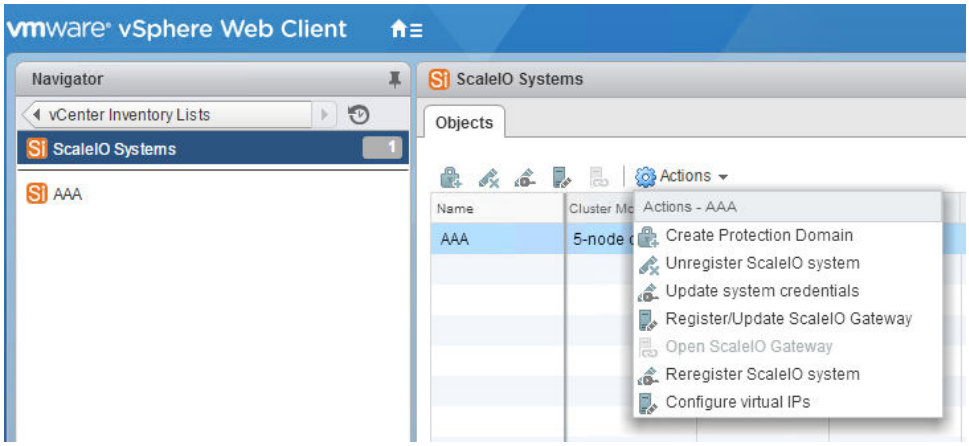

Note

You can also click the action icons in the menu or right-click the item to choose options from a list.

The **Create Protection Domain** dialog box appears:

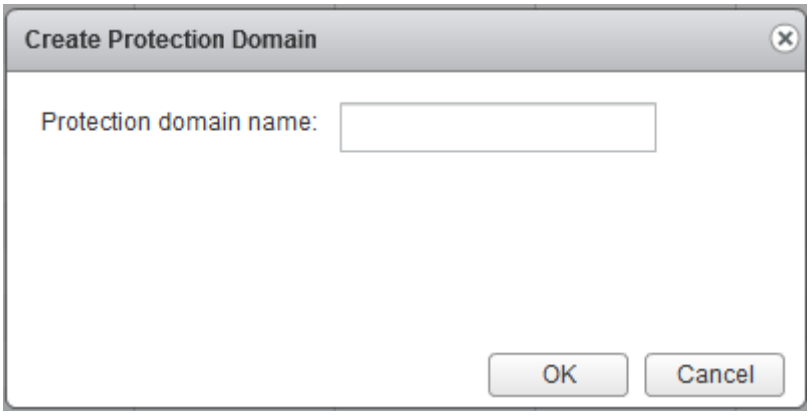

2. Enter a name for the Protection Domain, then click **OK**.

   The process is similar for the rest of the basic activities.

Note

If you intend to enable zero padding on a Storage Pool, you must do so before you add any devices to the Storage Pool. For more information, see Storage Pools on page 37, as well as the Command Quick Reference on page 99.

Configuring components—advanced

This section describes how to use the ScaleIO vSphere plug-in to perform activities that require a little more attention.

Registering an existing system

Register an existing ScaleIO system.

Procedure

1. From the main plug-in window, click **Register ScaleIO system**.
2. Enter the following information, then click **OK**:
   a. **Master MDM IP**: The IP address of the existing system's Master MDM
   b. **User name**: The username of the existing system
   c. **Password**: The password of the existing system

Creating, mapping, and unmapping volumes

This section describes how to use the plug-in to create, map, and unmap volumes in the VMware environment. You can map volumes to SDCs in the same step, or you can map the volume after it has been created.

Creating and mapping volumes

Volumes are created from devices in a Storage Pool.

Procedure

1. From the **Storage Pools** screen, click **Actions > Create volume**.

   The **Create Volume** dialog appears:

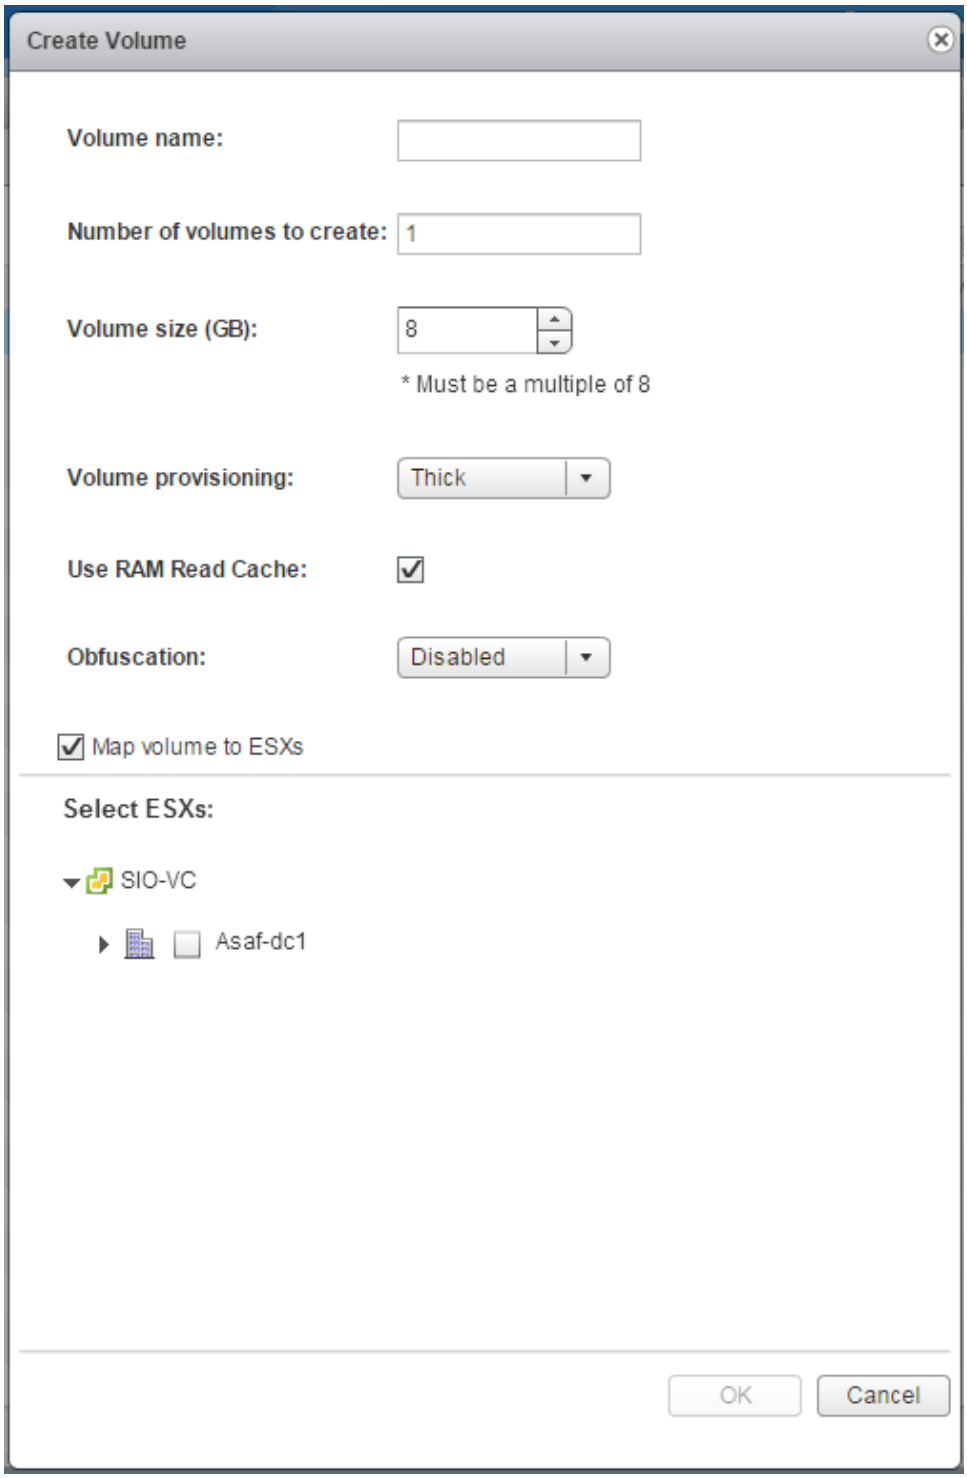

2. Enter the following information:
   - **Volume name**: Enter a name for the new volume.
   - **Number of volumes to create**: Enter the number of volumes to create. Multiple volumes appear as `volume_name-X`.
   - **Volume size**: Enter the size of the volume. This must be in multiples of 8 GB.
   - **Volume provisioning**: Select thick or thin provisioning.

- **Use RAM Read Cache**: Select to enable RAM Read Cache for the created volumes.Use of RAM Read Cache is determined by the policy for the Storage Pool and the volume.
- **Obfuscation** Select whether the volume should be obfuscated.

3. To map the volume to ESXs, perform the following steps:
   a. Select **Map volume to ESXs**.
   b. In the **Select ESXs** area, select the clusters or ESXs to which this volume should be mapped.
   c. Click **OK**.
   d. Enter the password for the ScaleIO administrative user.

   The following figure illustrates multiple volumes created:

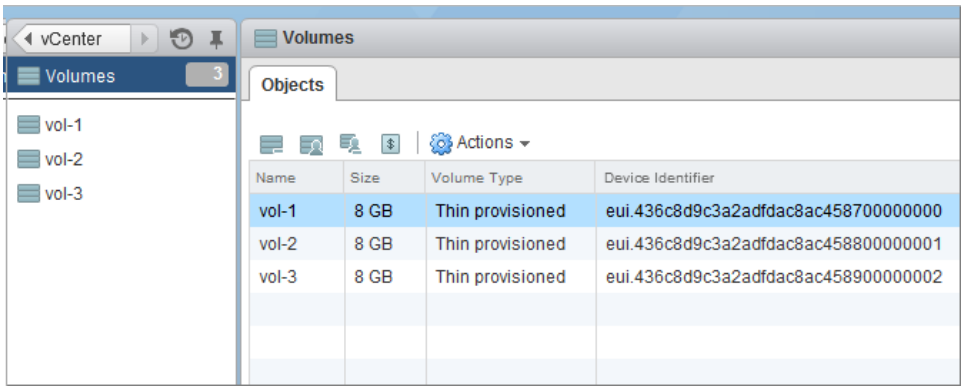

Mapping volumes

You can manually map volumes after they have been created, from the **Volumes** screen.

Procedure

1. From the **Volumes** screen, select a volume to map, then choose **Actions** > **Map a volume**.

   The **Map Volume to ESXs** dialog appears.

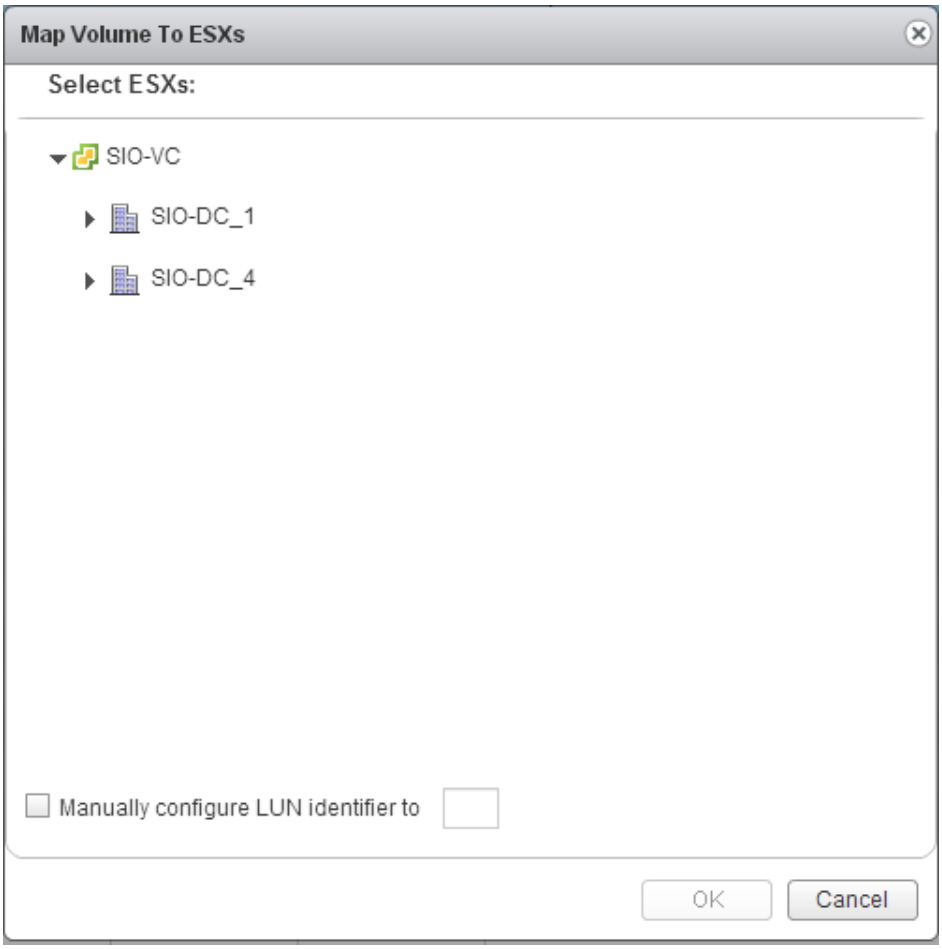

2. Select the clusters or ESXs to which this volume should be mapped.
3. To configure the LUN identifier manually, select **Manually configure LUN identifier to** and enter the identifier ID.
4. Click **OK**.

Unmapping a volume

This topic describes how to use the plug-in to unmap a volume from an ESX.

Procedure

1. From the **Volumes** screen, select the volume to unmap, and choose **Actions** > **Unmap volume**.

   The **Unmap Volume from ESXs** dialog appears.

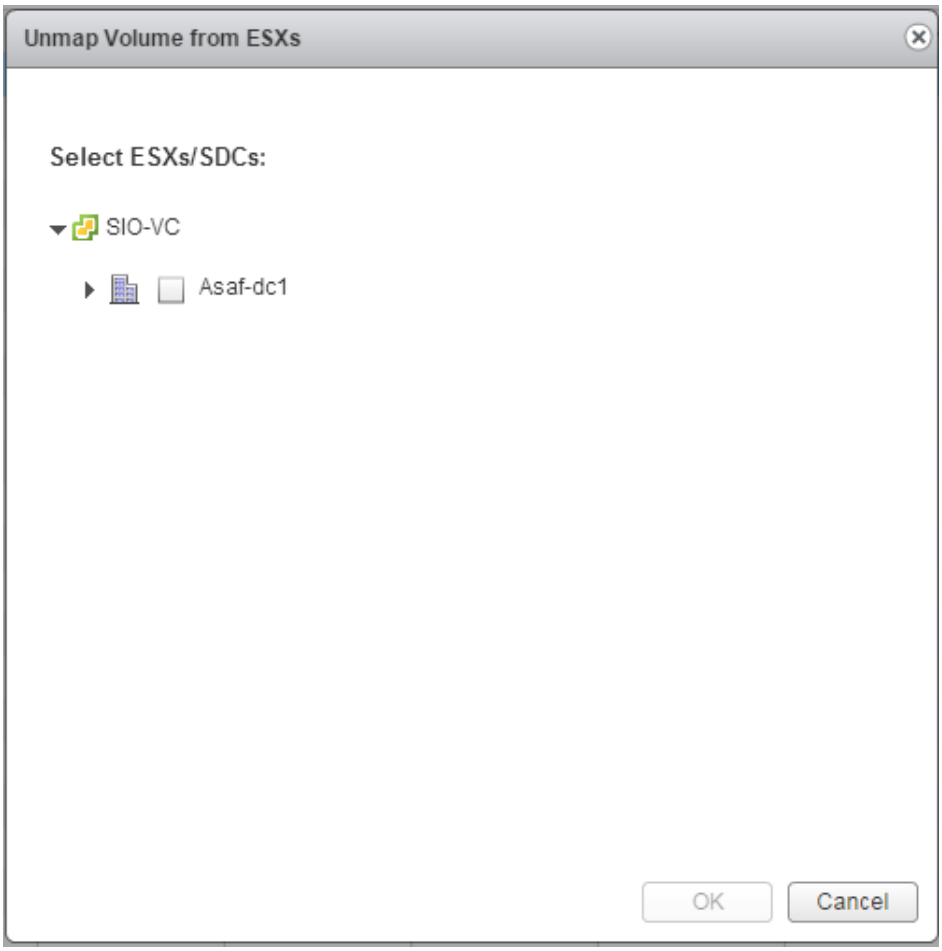

2. Select the ESXs or clusters from which to unmap the volume, then click **OK**.

Adding devices to an SDS

Use the vSphere plug-in to add devices to an SDS in a ScaleIO system.

In an RDM/VMDK-based ScaleIO system, you can add devices during and after the deployment. In a DirectPath-based system, you add devices only after the deployment.
You can add devices to a single SDS or to all SDSs in the system. The first option is quicker, but is limited to one SDS at a time.

All data on added devices will be erased.

Before using the vSphere plug-in to add devices, ensure that the devices have VD names. To add unnamed devices, use the SCLI `add_sds` command. For more information, see the "Add SDS device with vSphere plug-in fails in DirectPath environment" troubleshooting section in the *ScaleIO User Guide*.

Note

If you intend to enable zero padding on a Storage Pool, you must do so before you add any devices to the Storage Pool.

Procedure

1. From the SDSs screen of the ScaleIO vSphere plug-in, select one of the following:

- Right-click a specific SDS, then select **Add devices to a single SDS**.
- Right-click any SDS, then select **Add devices to ScaleIO system**.

The **Add Device** dialog appears. All devices that can be attached to the selected SDS are listed. For the system view, all SDSs are listed, and you can choose devices to add for each SDS. It may take a few moments to load the list of devices from the vCenter.

2. Add devices:
   - One-at-a-time:
     a. Select whether the device should be used for storage or to provide acceleration.
     b. Select the Storage Pool to which the devices should be assigned.
     c. To enable the use of devices that may have been part of a previous ScaleIO system. select **Allow the take over of devices with existing signature**.
     d. Click **OK**.
   - All devices on a server at once:
     a. Click **Select all devices**.
     b. Select whether to use the devices for storage or to provide acceleration.
     c. Select the Storage Pool to which the devices should be assigned.
     d. To enable the use of devices that may have been part of a previous ScaleIO system. select **Allow the take over of devices with existing signature**.
     e. Click **Assign**.
3. Confirm the action, by typing the ScaleIO password.
4. When the add operation is complete, click **Close**.

Results

The devices are added.

Upgrading an SDC

Upgrading an SDC is performed with the plug-in. This topic is described in the *ScaleIO Deployment Guide*.

Updating SDC parameters

Updating SDC parameters is necessary to ensure MDM-SDC communication when MDM IP addresses have been added or changed. This procedure, performed with the plug-in, is described in the *ScaleIO Deployment Guide*.

Note that there are differences in the way to perform this task, depending on the SDC version, as described in the guide.

Configuring virtual IP addresses

Configure virtual IP addresses in the vSphere Web plug-in.

Procedure

1. From the **ScaleIO Systems** screen, click **Actions** and select **Configure virtual IPs**.

2. In the **Configure virtual IPs** dialog box, select the network and enter a virtual IP address.

Figure 52 Configure virtual IPs dialog box

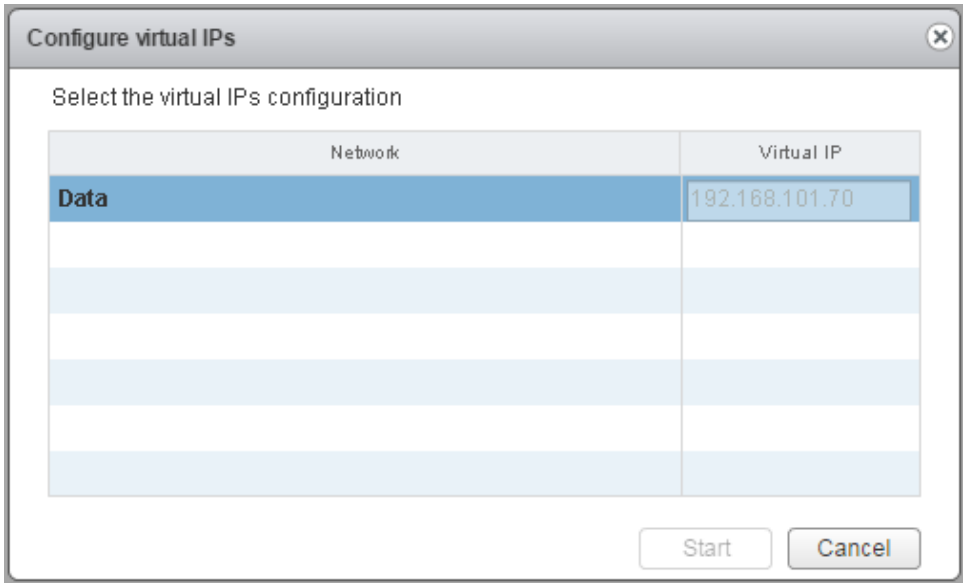

Note

Virtual IP addresses can only be added. To change or remove addresses, use the CLI.

After you finish

Update the SDC parameters to update the SDC configuration. For more information, see Updating SDC parameters on page 237.

Viewing components

To view an installed component, click it from the ScaleIO list.

Every component shows details that are relevant to the selected component.

For example, when you select **ScaleIO systems**, the following system details appear:

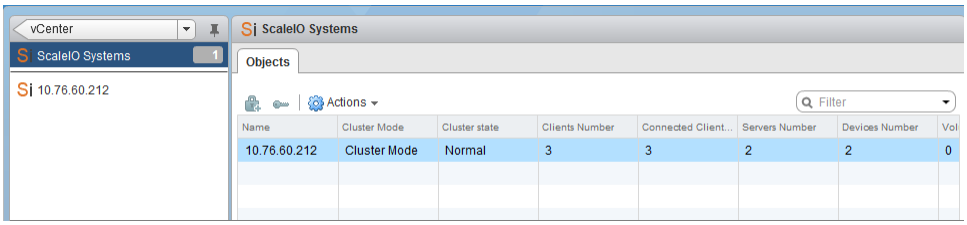

To drill down for more details, double-click the displayed details:

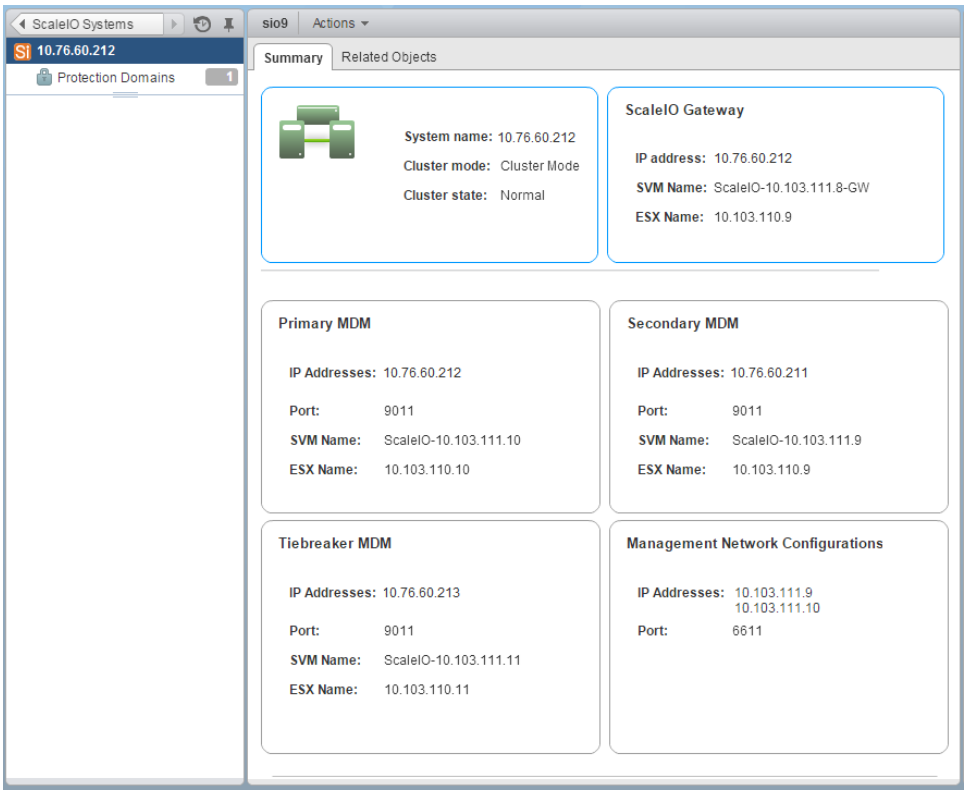

When you select **Storage Pools**, the following details appear:

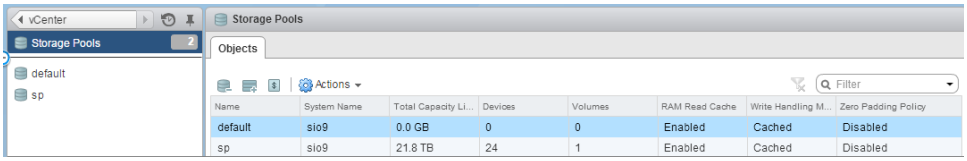

When you drill-down on this screen, the following details appear:

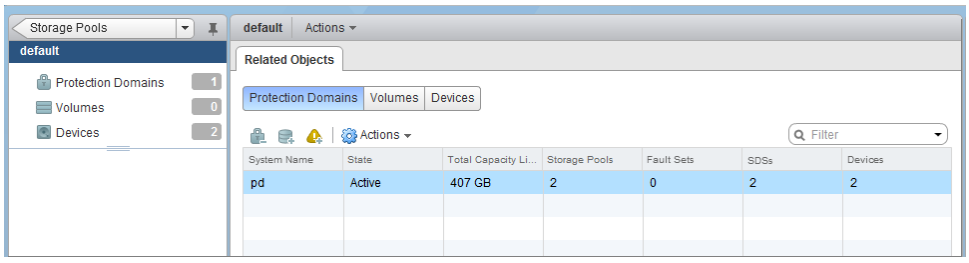

You can view the properties of all the ScaleIO components in the menu:

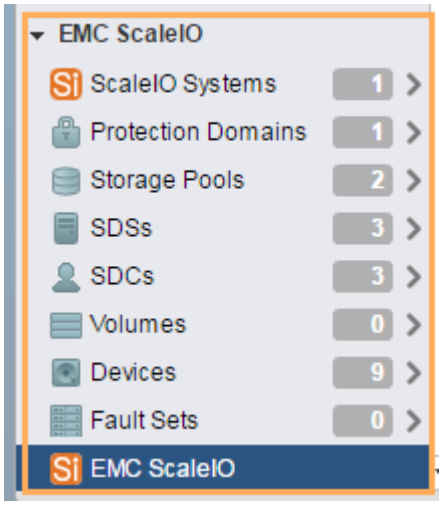
EMC ScaleIO
ScaleIO Systems 1
Protection Domains 1
Storage Pools 2
SDSs 3
SDCs 3
Volumes 0
Devices 9
Fault Sets 0
EMC ScaleIO

PART 4

Reference

The chapters in this part of the guide describe various topics related to advanced management. Chapters include:

Chapter 14, "Troubleshooting"

Chapter 15, "Frequently Asked Questions"

Chapter 16, "System events"

Chapter 17, "Return messages"

Chapter 18, " ScaleIO on Xen"

Chapter 19, "Configuring ScaleIO in OpenStack Environments"

Chapter 20, "REST API Reference"

Chapter 21, "SNMP trap support"

CHAPTER 14

Troubleshooting

This chapter describes ScaleIO troubleshooting. Topics include:

- Troubleshooting ScaleIO 244
- High latency encountered when S.M.A.R.T. hardware monitoring feature is enabled 244
- Adding a cache device to RFcache Storage Pool 244
- Add SDS device with vSphere plug-in fails in DirectPath environment 247
- After SDC installation, the ScaleIO SVM does not start automatically 247
- Application server does not see a ScaleIO volume 248
- Cannot log in to the Installation Manager after upgrade 248
- Certificate error when installing SDC on Windows servers 249
- VMware deployment failures 249
- Enabling acceleration in a new node 250
- S.M.A.R.T. hardware alerts are not displayed 253
- Gateway server recovery 253
- Installation Manager returns an error 257
- Installation with the Installation Manager fails 257
- LIA upgrade fails in Windows 2008 257
- Error when unmapping volume from ESXi server 258
- Older SDCs cannot communicate with newer SDSs after upgrade to v2.x 258
- Removing RFcache from Windows servers 259
- Removing RFcache leftovers from the Windows OS registry 259
- Replacing a faulty caching device in ScaleIO 259
- Replacing a faulty accelerated storage device in ScaleIO 261
- ScaleIO CLI or GUI cannot connect to an MDM 262
- ScaleIO Gateway fails to run 263
- SCLI add_sds command fails due to communication error or MDM going offline 263
- SVM manual memory allocation 263
- Viewing the status of volumes by drv_cfg --rescan command 265
- Virtual IP feature is not functional 266
- The VMware plug-in responds slowly 269
- Solving ScaleIO performance issues 270
- Mismatch in IO counters 270
- Deploying SVM on a node with a management IP address from a different subnet than the node with the SVM template 271
- Speeding up rebuild and rebalance processes 271
- RFcache is not installed on an upgraded Windows server 271
- SSD devices are not recognized in the operating system 272
- ESX competing thread setting resets on server reboot 273

Troubleshooting ScaleIO

This chapter describes solutions to issues that may arise. Most of the issues requiring troubleshooting are related to events and CLI return messages:

Event notifications
Upon changes or events in the system, ScaleIO may generate a system event that will be logged in a file. When addressing an event, look for the proper entry in the System Events appendix of the *ScaleIO User Guide*, according to the name of the event received. For each event that requires attention, the entry will include a description of a possible action to take. Follow the instructions before contacting EMC Support.

CLI messages
When issuing a CLI command, ScaleIO generates a textual response describing the outcome of the command. In some cases, the result may be a failure. Upon receiving a failure message, you can address it by looking for the proper entry in the Return Messages appendix of the *ScaleIO User Guide*, according to the text of the message received.

For each return message that requires attention, the entry will include a description of a possible action to take. Follow the instructions before contacting EMC Support.

High latency encountered when S.M.A.R.T. hardware monitoring feature is enabled

When the S.M.A.R.T. hardware monitoring feature is enabled, low-latency Storage Pools may exhibit higher-than-expected response times.

Disable both the watchdog and S.M.A.R.T. features on every SDS contributing to the Storage Pool in question.

Procedure

1. Add the following lines to the `conf.txt` file of the SDS:

```
tgt_dev__enable_metadata_polling = 0
```

2. Enter the SDS into maintenance mode.
3. Restart the SDS process:

```
pkill sds
```

4. Remove the SDS from maintenance mode.
5. Repeat this process for every SDS contributing to the Storage Pool.

Adding a cache device to RFcache Storage Pool

Add a cache device to an RFcache Storage Pool in a ScaleIO system.

Before you begin

Ensure that you have admin-level username and password for accessing ScaleIO.

Each RFcache pool on an SDS supports up to eight SDS devices as members in the RFcache Storage Pool. You can add additional devices while caching is in progress.

Procedure

1. Log in to an SDS.
2. Identify the SSD and HDDs to be accelerated and used as the cache devices, by running the following system commands:

SDS Operating System	Description
ESXi	Run the esxcli commands or open the vSphere View GUI to check the device properties.
Linux	Run the following commands: a. Run the command `fdisk -l` and check the size of the device in the output. b. Run the command `smartctl` and identify the type of device in the output.
Windows	Run either of these commands: • `wmic diskdrive list brief /format:list` and check the device properties in the output. • `wmic diskdrive get name,size,model` and check the device name, size, and model number in the output.

3. Log in to the Master MDM.
4. Add an SSD device to the default cache pool, by running the following command:

```
scli --add_sds_rfcache_device (--sds_id <ID> |--sds_name
<NAME> | --sds_ip <IP> [--sds_port <PORT>])
--rfcache_device_path <PATH> [--rfcache_device_name <NAME>]
```

Example:

```
scli --add_sds_rfcache_device --sds_name SDS1 --
rfcache_device_path /dev/sdb --rfcache_device_name
rfcache1_SDS1_sdb
```

5. Repeat the previous steps, from the beginning of the procedure, for each new SDS device to be configured in the RFcache Storage Pool.
6. Verify the status of acceleration per SDS, by running the following command:

```
scli --query_sds (--sds_id <ID> | --sds_name <NAME> |--sds_ip
<IP>)
```

Example:

```
scli --query_sds --sds_ip 192.168.1.6
```

Output, similar to the following, appears, for different command types:

- Command to verify the RFcache status:

  ```
  scli --query_sds (--sds_id <ID> | --sds_name <NAME> |--
  sds_ip <IP>)
  ```

  Example:

  ```
  scli --query_sds --sds_ip 9.10.21.141
  ```

  Output:

  ```
  SDS 4eecc17300000002 Name: SDS_10.136.215.203 Version:
  2.0.13000
  Protection Domain: 928cd6e100000000, Name:
  Default_Protection_Domain
  DRL mode: Volatile
  Authentication error: None
  IP information (total 2 IPs):
           1: 9.10.21.141      Role: All (SDS and SDC)
           2: 9.10.121.141     Role: All (SDS and SDC)
          Port: 7072
  RAM Read Cache information:
          128.0 MB (131072 KB) total size
          Cache is enabled
          RAM Read Cache memory allocation state is PENDING
  Rfcache enabled
  .
  .
  .

  Rfcache device information (total 1 devices):
           1: Name: 50000396CC883494  Path: /dev/sdc  Original-
  path: /dev/sdc  Size 762495MB  ID: e177c11700020000
  ```

- Command to verify Storage Pool status:

  ```
  scli --query_storage_pool (((--protection_domain_id <ID> |
  --protection_domain_name <NAME>) --storage_pool_name
  <NAME>) | --storage_pool_id <ID>)
  ```

  Example:

  ```
  scli --query_storage_pool --storage_pool_name
  Cached_Hdd_Pool_1 --protection_domain_name
  Default_Protection_Domain
  ```

  Output:

  ```
  Storage Pool Cached_Hdd_Pool_1 (Id: ada1bb8500000000) has 0
  volumes and 23.0 TB (23536 GB) available for volume allocation
          The number of parallel rebuild/rebalance jobs: 2
          Rebuild is enabled and using Limit-Concurrent-IO
  ```

```
policy with the following parameters:
        Number of concurrent IOs per device: 1
        Rebalance is enabled and using Favor-Application-IO
policy with the following parameters:
        Number of concurrent IOs per device: 1, Bandwidth
limit per device: 10240 KB per second
        Background device scanner: Disabled
        Zero padding is enabled
        Spare policy: 36% out of total
        Checksum mode: disabled
        Doesn't use RAM Read Cache
        Uses Flash Read Cache
```

Results

The new cache devices are configured in the RFcache Storage Pool in ScaleIO.

Add SDS device with vSphere plug-in fails in DirectPath environment

When attempting to add an SDS device with the vSphere plug-in, an error is displayed. Devices to be added in a DirectPath environment must have a Virtual Device (VD) name before being added.

When a ScaleIO system is configured to use ESXi DirectPath, adding an SDS device with the ScaleIO vSphere plug-in will fail if the device corresponds to a PERC H730P VD, that was previously created with the name field not set.

The result is that the device is not added and an error message, similar to the following, is displayed: "slot_0_devNum_0_hdd (N/A) - Failed: Add storage device ScaleIO-1d7c870d to Storage Pool SP2 (ScaleIO - SDS device not found)".

Workaround:

Before adding a VD storage device as an SDS device, ensure that the VD has a name. If a name was not specified when the device was created, use PERCCLI to assign a name before adding. A reboot of the SDS and the SVM may be required. (You can also change the name using the PERC H730P configuration management utilities by accessing System Setup during boot.)

An alternate way of adding a device (even without a named VD) is using the SCLI `add_sds_device` command. For example:

```
scli --add_sds_device --sds_id 107583d600000001 --storage_pool_name
SP2 --device_path /dev/sdb
```

After SDC installation, the ScaleIO SVM does not start automatically

After deployment is completed, set all ScaleIO VMs to start automatically:

Procedure

1. Click the ESX **Configuration** tab.
2. From the **Software** section, click **Virtual Machine Startup/Shutdown**.
3. Click **Properties**.
4. In the dialog box, select **Allow virtual machines to start and stop automatically with the system**.

5. Select the SVM and move it to the **Automatic Startup** list.
6. Click **OK**.
7. Repeat this process for all SVMs.

Application server does not see a ScaleIO volume

Perform the following steps:

Procedure

1. Check if the ScaleIO system is operational:

```
scli --mdm_ip <mdm IP> --query_all
```

2. Check if the volume is mapped to any of the SDC servers by running the `scli --query_all_volumes` command.
3. Determine if the SDC is installed on the server:
   - Linux: Run `rpm -qa | grep sdc`
   - ESX: Run `esxcli software vib list|grep sdc`
4. Determine if the SDC is connected to an MDM:

```
scli --mdm_ip <mdm IP> --query_all_sdc
```

5. Ensure that the MDM management IP address is up and running.
6. On an application server, rescan for new volumes:
   - Linux: Run `/opt/emc/scaleio/sdc/bin/drv_cfg --rescan`
   - Windows: Run `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg --rescan`
   - ESX: Rescan for new devices

Cannot log in to the Installation Manager after upgrade

If you upgrade the IM, and are then unable to log in, restart the EMC ScaleIO Gateway service.

Certificate error when installing SDC on Windows servers

Due to a Microsoft Windows limitation, when installing an SDC component on a Windows 2008 R2 server, a certificate error, similar to the following may be displayed:

To solve this issue, ensure that Microsoft Security Update KB3033929 is installed.

VMware deployment failures

Correct failures during a 5-node deployment with the ScaleIO vSphere plug-in deployment wizard.

The deployment may fail with any of the following errors:

- Failed: Install ScaleIO SDS module (Cannot connect to MDM cluster)
- Failed: install ScaleIO RFCACHE module (Cannot connect to MDM cluster)
- Failed: Configure SDC driver on ESX (Cannot connect to MDM cluster)

To fix these issues, click **Continue deployment** in the wizard. Deployment should continue successfully.

Enabling acceleration in a new node

Enable acceleration in a new SDS node in a ScaleIO system.

Before you begin

Ensure that you have:

- Admin-level username and password for accessing ScaleIO
- The RFcache package (`xcache`) for installation on a new SDS
- This task requires other ScaleIO procedures, that are described in other documents:
  - Installing an SDS in ScaleIO
  - Adding an SDS to the MDM
  - Enabling RFcache, using the Installation Manager (IM)
  - Accessing the SVM

Procedure

1. Install an SDS in the ScaleIO system and add it to the MDM.
2. Enable RFcache on the SDS:

SDS OS	Description
ESXi	Access the SVM and then run the following command: `rpm -i EMC-ScaleIO-xcache-2.0-14000.0.sles11.3.*SVM*.x86_64.rpm`
Linux	Enable RFcache using the Installation Manager.
Windows	Enable RFcache using the Installation Manager.

3. Identify the SSD and HDDs to be accelerated and used as cache devices:

SDS OS	Description
ESXi	Run the esxcli commands or open the vSphere View GUI to check the device properties.
Linux	Run the following commands: a. Run the command `fdisk -l` and check the size of the device in the output. b. Run the command `smartctl` and identify the type of device in the output.
Windows	Run either of the following commands: • `wmic diskdrive list brief /format:list` and check the device properties in the output. • `wmic diskdrive get name,size,model` and check the device name, size, and model number in the output.

4. Log in to the Master MDM.

5. Add an SSD device to the default cache pool:

```
scli --add_sds_rfcache_device (--sds_id <ID> |--sds_name
<NAME> | --sds_ip <IP> [--sds_port <PORT>])
--rfcache_device_path <PATH> [--rfcache_device_name <NAME>]
```

Example:

```
scli --add_sds_rfcache_device --sds_name SDS1 --
rfcache_device_path /dev/sdb --rfcache_device_name
rfcache1_SDS1_sdb
```

6. Repeat the previous steps, from the beginning of the procedure, for each new SDS device to be configured in the RFcache Storage Pool.

7. Enable caching on the Storage Pool in the Protection Domain:

```
scli --set_rfcache_usage (((--protection_domain_id <ID> | --
protection_domain_name <NAME>) --storage_pool_name <NAME>) |
--storage_pool_id <ID>) [--use_rfcache | --dont_use_rfcache]
```

Example:

```
scli --set_rfcache_usage --protection_domain_name PD1 --
storage_pool_name SP1 --use_rfcache
```

8. Verify the status of acceleration per SDS:

```
scli --query_sds (--sds_id <ID> | --sds_name <NAME> |--sds_ip
<IP>)
```

Example:

```
scli --query_sds --sds_ip 192.168.1.6
```

Output, similar to the following, appears, for different command types:

- Command to verify the RFcache status:

```
scli --query_sds (--sds_id <ID> | --sds_name <NAME> |--
sds_ip <IP>)
```

Example:

```
scli --query_sds --sds_ip 9.10.21.141
```

Output:

```
SDS 4eecc17300000002 Name: SDS_10.136.215.203 Version:
2.0.13000
Protection Domain: 928cd6e100000000, Name:
Default_Protection_Domain
DRL mode: Volatile
Authentication error: None
IP information (total 2 IPs):
         1: 9.10.21.141      Role: All (SDS and SDC)
         2: 9.10.121.141     Role: All (SDS and SDC)
        Port: 7072
RAM Read Cache information:
        128.0 MB (131072 KB) total size
        Cache is enabled
        RAM Read Cache memory allocation state is PENDING
Rfcache enabled
.
.
.

Rfcache device information (total 1 devices):
         1: Name: 50000396CC883494  Path: /dev/sdc  Original-
path: /dev/sdc  Size 762495MB  ID: e177c11700020000
```

- Command to verify Storage Pool status:

```
scli --query_storage_pool (((--protection_domain_id <ID> |
--protection_domain_name <NAME>) --storage_pool_name
<NAME>) | --storage_pool_id <ID>)
```

Example:

```
scli --query_storage_pool --storage_pool_name
Cached_Hdd_Pool_1 --protection_domain_name
Default_Protection_Domain
```

Output:

```
Storage Pool Cached_Hdd_Pool_1 (Id: ada1bb8500000000) has 0
volumes and 23.0 TB (23536 GB) available for volume allocation
        The number of parallel rebuild/rebalance jobs: 2
        Rebuild is enabled and using Limit-Concurrent-IO
policy with the following parameters:
        Number of concurrent IOs per device: 1
        Rebalance is enabled and using Favor-Application-IO
policy with the following parameters:
        Number of concurrent IOs per device: 1, Bandwidth
limit per device: 10240 KB per second
        Background device scanner: Disabled
        Zero padding is enabled
        Spare policy: 36% out of total
        Checksum mode: disabled
        Doesn't use RAM Read Cache
        Uses Flash Read Cache
```

9. Repeat this entire procedure on the rest of the SDS nodes.

Results

The new nodes are configured in the RFcache Storage Pool in ScaleIO.

S.M.A.R.T. hardware alerts are not displayed

Hardware alerts are not displayed in the ScaleIO GUI.

1. Ensure that the proper driver/disk utility is installed on the servers:
   - LSI: storcli
   - HP servers: hpssacli
   - Dell servers: perccli

   The CLI utility enables you to display the hardware alerts for the RAID controller and disk status.
2. Ensure that the LIA process is running on the server.
   LIA is used for alert display.
3. Ensure that the specific RAID controller is supported on the server.

Gateway server recovery

The following procedures describe how to recover a ScaleIO gateway when the physical host server fails:

Recovering Gateway on a new server during upgrade

Recover the ScaleIO Gateway (GW) on a new server and continue with the process of ScaleIO upgrade, when the GW server fails during ScaleIO upgrade.

If the GW server fails during an upgrade, the Installation Manager (IM) will also go down. You must recover the GW first, and then the IM. This procedure will cause the MDM to use the new GW in place of the failed one.
To recover the Gateway on a new server, perform the following steps:

Procedure

1. Install the Gateway on a new server (identical to the one that crashed), using any of the following options:
   - Linux OS: Use the GW RPM file.
   - Windows OS: Use the GW MSI file.
2. If the LIA-trusted IPs feature was enabled during installation, you must perform the steps described in Recovering IM with LIA trusted IP feature enabled on page 255 before continuing with this procedure.
3. Open the `gatewayUser.properties` file, located at `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/gatewayUser.properties`, and edit the IM parameters.

Parameter	Description	Action
`upgrade.mdm.data.ips`	Configures MDM data IPs	Enter comma-separated IP addresses for different cluster nodes. • If a single node has multiple IP addresses, then enter plus-sign-separated IPs. • Ensure that the sequence of items in the list matches the order of data, management (mgmt), and role for the same node.

Parameter	Description	Action
		• Reserve the empty values for Tie Breaker management (TB mgmt) IP addresses. Example: • `upgrade.mdm.data.ips` = 192.100.1.10,192.100.1.49,192.100.1.9,192.100.1.17,192.100.1.11 • `upgrade.mdm.mgmt.ips` = 10.76.60.41,10.76.60.48,10.76.60.41,,
`upgrade.mdm.mgmt.ips`	Configures MDM management IPs	Enter comma-separated IP addresses for different cluster nodes. • If a single node has multiple IP addresses, then enter plus-sign-separated IP addresses • Ensure that the sequence of items in the list matches the order of data, management (mgmt), and role for the same node • Reserve the empty values for Tie Breaker management (TB mgmt) IP addresses. Example: • `upgrade.mdm.data.ips` = 192.100.1.10,192.100.1.49,192.100.1.9,192.100.1.17,192.100.1.11 • `upgrade.mdm.mgmt.ips` = 10.76.60.41,10.76.60.48,10.76.60.41,,
`upgrade.mdm.role`	Configures MDM roles	Enter the role of the MDM server; mdm, slave, or tb (Tie Breaker). The line of the MDM role should correlate to the MDM data IP and MDM management IP addresses. Example: `upgrade.mdm.role` = mdm,slave,slave,tb,tb
`upgrade.mdm.version.target`	Configures the target version of MDM	Enter the MDM target version to which you want to upgrade. The first item in the `version.target` must represent the first cluster node. Example: `upgrade.mdm.version.target` = 2.0-828.0,2.0-828.0,2.0-828.0
`upgrade.mdm.version.orig`	Configures the original version of MDMs	Enter the MDM's original version, from which you are upgrading. The first item in the `version.orig` must represent the first cluster node. Example: `upgrade.mdm.version.orig` = 1.32-255.0,1.32-255.0,1.32-255.0
`upgrade.mdm.actor.port`	Configures MDM actor ports	Enter comma-separated MDM actor ports, as an ordered list. Each entry is a relevant option, rollback or upgrade. Option selection relates to the operation, version upgrade or rollback that was being performed before the GW crashed. Example: `upgrade.mdm.status` = upgrade
`upgrade.sds.mm_list`	Removes the SDS from Maintenance mode	Change the status of the SDS, which was in maintenance mode before the GW server crash using the `upgrade.sds.mm_list` parameter. Example:

Parameter	Description	Action
		`upgrade.sds.mm_list` = 10.76.60.32,10.76.60.33,10.76.60.34,10.76.60.35,10.76.60.36

Results

The IM parameters in the `gatewayUser.properties` file match with the scenario that existed before the GW server crash.

Recovering IM with LIA trusted IP feature enabled

Describes how to recover the Installation Manager (IM) when the LIA-trusted IP feature was enabled in the crashed IM.

When the IM crashes during system upgrade and the LIA trusted IPs feature was enabled with a faulty IM, you are required to update the LIA's white list with the new IM IPs.

Note

When this procedure is required and the Gateway needs to be recovered, you must perform this procedure before recovering the Gateway.

To update the LIA's white list, perform the following steps on all nodes:

Procedure

1. Access the LIA node in the system.
2. Enter the IP addresses:

OS	Steps

Linux	a. Go to: `/opt/emc/scaleio/lia/bin/`
	b. Run (on one line):
	`./update_conf_bin 7 lia_trusted_ips <Comma_separated_Scaleio_gw_IPs> 1`
	Ensure that you enter all of the ScaleIO GW (IM) hosts' available IP addresses, even if those addresses are not used by ScaleIO.
	Example:
	`./update_conf_bin 7 lia_trusted_ips 1.2.3.4,5.6.7.8 1`

Windows	a. Go to: `C:\Program Files\EMC\scaleio\sds\bin` b. Run (on one line): `update_conf.exe 7 lia_trusted_ips <Comma_separated_Scaleio_gw_IPs> 1` Ensure that you enter all of the ScaleIO GW (IM) hosts' available IP addresses, even if those addresses are not used by ScaleIO.

OS	Steps
	Example: `update_conf.exe 7 lia_trusted_ips 1.2.3.4,5.6.7.8 1`

SNMP configuration recovery post gateway crash during upgrade

Recover SNMP configuration when the gateway crashes during ScaleIO upgrade.

There are two methods of recovering SNMP configuration:

- Recovering SNMP configuration using Installation Manager
- Manually recovering SNMP configuration

Recovering SNMP configuration using Installation Manager

You can use the Installation Manager to recover the SNMP configuration, when the Gateway crashes during a ScaleIO upgrade.

Procedure

1. Delete the alertservice user from the MDM:
   a. Log in to the MDM with the admin user.

      Example:

      ```
      scli --login --username admin --password Scaleio123
      ```

   b. List all existing users:

      ```
      scli --query_users
      ```

   c. Delete the designated user:

      ```
      scli -delete_user (-user_id <ID> | --username <NAME>)
      ```

      Example:

      ```
      scli --delete_user --username alertservice
      ```

   d. List all existing users, and ensure that the deleted user is not listed:

      ```
      scli --query_users
      ```

2. Open the IM GUI and extend the ScaleIO system, using an updated CSV file.
3. In the **Installation Configuration** window, select **Set advanced options (optional)**.

   The **Advanced Configuration** section is displayed.

4. In the **Config. Options** section, select **Enable alert service (required for SNMP and ESRS)**.
5. In the **traps Receiver IP** field, type the traps receiver IP address.
6. Complete the extend operation.

Manually recovering SNMP configuration

Recover SNMP configuration manually, when the gateway crashes during ScaleIO upgrade.

Procedure

1. Create a Lockbox.
2. On the Gateway, edit the `gatewayUser.properties` file:

   For more information, see the *EMC ScaleIO User Guide*

   a. Change the value of `features.enable_snmp` to *true*.

   b. Enter IP addresses in `snmp.traps_receiver_ip`.

      ScaleIO supports up to two comma-separated TRAP receivers.

   c. In `snmp.sampling_frequency`, update the MDM sampling period (optional step).

   d. In `snmp.resend_frequency`, update the time period to re-send the already existing TRAPs (optional step).
3. Restart the scaleio-gateway service.

Installation Manager returns an error

In the Installation Manager, if an `HTTP 404 Webpage not available` error (or similar) is displayed, the IM may have been disabled.

For more information, refer to the "Configuring the Installation Manager" section in the *Deployment Guide*.

Installation with the Installation Manager fails

If the following error message appears when installing with the Installation Manager,

```
Command failed: Could not connect to <IP_address>
```

It could be that this node is not in an accessible network.

LIA upgrade fails in Windows 2008

When upgrading or installing the ScaleIO system on Windows 2008R2-based servers, the upgrade or installation may fail and generate this error message: `<package X> for OS windows is not signed`.

Due to an issue in Windows 2008, the certificate issuer might not be recognized. This will result in failure to upgrade or install the LIA MSI in Windows 2008.

To resolve the problem, use the following workaround:

Procedure

1. Download and extract the `roots.zip` root certificates package, located at https://www.symantec.com/theme/roots.
2. Manually install all the certificates in the VeriSign Root Certificates subfolders:
   a. Open a subfolder.
   b. Double-click a certificate, and select **Install Certificate**.
   c. Select **Place all the certificates in the following store**, and then browse to and select the **Trusted Root Certification Authorities** store.
   d. Click **OK**, then **Next**, and then **Finish** .
   e. Review and approve the security warning.

      A message appears, informing you that the import was successful.

   f. Repeat for each certificate.
3. Retry the upgrade or installation.

   The upgrade or installation should now succeed.

Error when unmapping volume from ESXi server

Correct failure that occurs when unmapping volumes from an ESXi server.

Before you begin

You need the login credentials for the ESXi host where the unmapping is not succeeding.

Sometimes, when unmapping volumes from an ESXi server, a " Could not detach LUN" error is displayed. Use the following steps to unmap the volume.

Procedure

1. Use SSH to log in to the ESXi host.
2. Run this command:

```
esxcli storage core device detached remove --all
```

3. From ScaleIO, unmap the volume again.

Results

The unmap command succeeds.

Older SDCs cannot communicate with newer SDSs after upgrade to v2.x

After upgrading to v2.x, older SDCs may not be able to communicate with newer SDSs.

If the IP addresses of a v2.x SDS are defined in IPv6 only, older-generation SDCs using IPv4 will not be able to communicate with the SDS. This can cause I/O errors. This can occur when v1.32 SDCs are added to a v2.0 system, as v2.0 supports IPv6, whereas earlier versions do not.

To resolve the problem, if the SDCs are to remain pre-v2.0:

- Retain at least one IPv4 address on the SDS.
- Before removing all the IPv4 IP addresses from the SDSs, ensure that all SDCs are upgraded to v2.x and then rebooted .

Removing RFcache from Windows servers

To remove RFcache from Windows servers, use the Windows Control Panel Add/Remove function to remove xcache.

If SDC is also installed on that server, you must remove SDC first, restart the machine, and then remove xcache.

Removing RFcache leftovers from the Windows OS registry

Run a batch file to remove RFcache leftovers from the Windows registry.

Uninstalling RFcache leaves some leftover entries in the Windows registry. This will cause errors the next time RFcache is installed. To remove the leftover entries in the registry, use the `Clean_XC_registry.bat` file included in the ScaleIO Windows download.

Procedure

1. Locate the `Clean_XC_registry.bat` file that is included in the ScaleIO Windows download.
2. From the Windows command prompt, run the `Clean_XC_registry.bat` file.

Replacing a faulty caching device in ScaleIO

Replace a faulty device that provides acceleration (caching) in a ScaleIO system.

Before you begin

Ensure that you have:

- Administrator authentication credentials to run ScaleIO SCLI commands
- Access to the ScaleIO nodes
- Any of the following parameters:
  - SDS name
  - SDS IP address
  - SDS ID

This procedure applies to ScaleIO software-only systems. It does not apply to AMS-supported ScaleIO systems.

Procedure

1. Log in to the scli interface.
2. Identify the failed RFcache device:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Output, similar to the following, appears for the RFcache devices:

```
Rfcache device information (total 1 devices):
         1: Name: N/A  Path: /dev/sdi  Original-path: /dev/sdi
Size 952719MB  ID: dbfb70fc00000000
```

For a faulty device, the command output displays the fault (error) at the end of the output, while the output for a healthy device does not.

3. Remove the faulty RFcache device from the Storage Pool:

```
scli --remove_sds_rfcache_device (--rfcache_device_id <ID> |
((--sds_id <ID> | --sds_name <NAME> | --sds_ip <IP> [--
sds_port <PORT>]) (--rfcache_device_name <NAME> | --
rfcache_device_path <PATH>)))
```

Example:

```
scli --remove_sds_rfcache_device --rfcache_device_id
7cf1ce1f00000000
```

4. Physically replace the faulty device, using the relevant system and vendor guidelines.
5. Add a new device:

```
scli --add_sds_rfcache_device (--sds_id <ID> |--sds_name
<NAME> | --sds_ip <IP> [--sds_port <PORT>])
--rfcache_device_path <PATH> [--rfcache_device_name <NAME>]
```

Example:

```
scli --add_sds_rfcache_device --sds_name SDS1 --
rfcache_device_path /dev/sdb --rfcache_device_name
rfcache1_SDS1_sdb
```

6. Reset the device error counter:

```
scli --clear_sds_rfcache_error (--sds_id <ID> |--sds_name
<NAME> | --sds_ip <IP> [--sds_port <PORT>])
```

7. Verify the status of acceleration per SDS:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Example:

```
scli --query_sds --sds_name SDS1
```

The command output should display the device state without errors (that is, the replaced device in healthy state); otherwise, repeat the procedure.

Results

The faulty caching device is replaced in the ScaleIO system.

Replacing a faulty accelerated storage device in ScaleIO

Replace a faulty accelerated (cached) storage device in a ScaleIO system.

Before you begin

Ensure that you have:

- Administrator authentication credentials to run ScaleIO SCLI commands
- Access to the ScaleIO nodes
- Any one of the following parameters:
  - SDS name
  - SDS IP address
  - SDS ID

This procedure applies to ScaleIO software-only systems. It does not apply to AMS-supported ScaleIO systems.

Procedure

1. Log in to the scli interface.
2. Identify the failed cached storage HDD device:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Output, similar to the following, appears:

```
Device information (total 7 devices):
         1: Name: N/A  Path: /dev/sdb  Original-path: /dev/sdb
ID: ffcfa55700000000
                Storage Pool: sp1, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         2: Name: N/A  Path: /dev/sdc  Original-path: /dev/sdc
ID: ffcfa55800000001
                Storage Pool: sp1, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         3: Name: N/A  Path: /dev/sdd  Original-path: /dev/sdd
ID: ffcfa55900000002
                Storage Pool: sp1, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         4: Name: N/A  Path: /dev/sde  Original-path: /dev/sde
ID: ffcfa55a00000003
                Storage Pool: sp2, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         5: Name: N/A  Path: /dev/sdf  Original-path: /dev/sdf
ID: ffcfa55b00000004
                Storage Pool: sp2, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         6: Name: N/A  Path: /dev/sdg  Original-path: /dev/sdg
ID: ffcfa55c00000005
                Storage Pool: sp2, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         7: Name: N/A  Path: /dev/sdh  Original-path: /dev/sdh
ID: ffcfa55d00000006
                Storage Pool: sp3, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
```

For a faulty device, the command output displays the fault (error) at the end of the output, while the output for a healthy device does not.

3. Remove the faulty device from the Storage Pool:

```
scli --remove_sds_device (--device_id <ID> |((--sds_id <ID> |
--sds_name <NAME> | --sds_ip <IP> [--sds_port <PORT>]) (--
device_name <NAME> | --device_path <PATH>)))
```

4. Physically replace the faulty device, using the relevant system and vendor guidelines.
5. Add a new device, as a cached device, to the Storage Pool:

```
scli --add_sds_device (--sds_id <ID> | --sds_name <NAME> | --
sds_ip <IP> [--sds_port <PORT>])
--device_path <PATH> [--device_name <NAME>] (--
storage_pool_name <NAME>) | --storage_pool_id <ID>)
```

When the device is added to the Storage Pool, the device starts being cached.

6. Verify the status of the replaced device:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Example:

```
scli --query_sds --sds_name SDS1
```

The command output should display the device state without errors (that is, the replaced device in healthy state); otherwise, repeat the procedure.

Results

The faulty cached device is replaced in the ScaleIO system.

ScaleIO CLI or GUI cannot connect to an MDM

Perform the following steps:

Procedure

1. Ping the MDM IP address to ensure you have connectivity.
2. Ensure that you are connecting to the IP address of the Master MDM.

   If the MDM ownership has changed, try to connect to the IP address of the Slave MDM.
3. Check if the MDM is running, by typing the following command:

   ```
   ps -ef| grep mdm
   ```
4. Ensure that the management IP address is up and running.

ScaleIO Gateway fails to run

The ScaleIO Gateway fails to run, or it runs but does not listen to any communications. In some cases, gateway installation fails.

Any of the above scenarios may occur due to gateway port collision. The ScaleIO Gateway requires listening ports 80 and 443 when running as root, or listening ports 8080 and 8443 when running as non-root.

These ports are required in order to receive communications from the OpenStack controller.

On Linux servers, to verify whether the required ports are already in use, run:

```
netstat -tupln
```

On Window servers, run:

```
netstat -nab
```

None of the required ports (80 or 443 for root, or 8080 or 8443 for non-root) should be included in the command output.

If any of the required ports is already in use by an application other than ScaleIO, the services currently using the ports should be shut down for the duration of the installation and use of the gateway. Alternatively, gateway installation should be delayed until the required ports are made available.

For more information, see Port usage and changing default ports on page 278.

SCLI add_sds command fails due to communication error or MDM going offline

When you run the SCLI `add_sds` command, the command may fail either due to a communication error or the Master MDM going offline.

Normally, when you run `add_sds`, the SDS receives a request and returns acknowledgment (ACK) to the MDM. However, after you initiate an `add_sds` command, if the MDM goes down for any reason, the MDM may not receive the ACK message (because the MDM is offline), but the SDS presumes that it is attached to the MDM.

Thus, when the MDM does not receive the ACK message, retrying the `add_sds` command fails with status as `SDS is already attached to this MDM`.

To resolve this issue, it is recommended that you clean the SDS configuration by running the `add_sds` command with `force_clean` flag and then re-run `add_sds`.

SVM manual memory allocation

When using the plug-in for a clean deployment, SVM memory allocation is performed automatically. In the following cases, SVM memory allocation must be performed manually:

- Manual deployment on VMware.

- Extending an existing SVM with a new ScaleIO role/component, whether this is being done with the plug-in or manually.
  Workaround: Perform all the parts of step 1 and step 2 before extending the additional role/component on the SVM. Perform the steps on one SVM at a time.
- Changing the SDS performance profile, post deployment, or after an upgrade from v1.32.x to v2.0.
  Workaround: Perform all the parts of step 1 one SVM at a time.
- Post backend upgrade from v1.32.x to v2.0.
  Workaround: Perform all the parts of step 1 and step 2 on all SVMs that were upgraded, in the following order, one SVM at a time.
- Adding capacity to an SDS that was deployed with partially-populated capacity.
  Workaround: Perform all the parts of step 1 one SVM at a time.

Procedure

1. For SVMs that are SDS-only, perform the following:
   a. Move the SDS to maintenance mode (MM).
   b. Shut down the SVM.
   c. Increase SVM memory, according to the formula below.
   d. Power up the SVM.
   e. Exit MM.
2. For SVMs that are MDM (Master, Slave, or TB, may contain SDS, also):
   a. Start with Slaves and TBs:
      a. Move the SDS to maintenance mode (MM).
      b. Shut down the SVM.
      c. Increase SVM memory, according to the formula below.
      d. Power up the SVM.
      e. Exit MM.
   b. Proceed with the Master MDM:
      a. Switch ownership, so the Master MDM is now a Slave MDM.
      b. Move the SDS to maintenance mode (MM).
      c. Shut down the SVM.
      d. Increase SVM memory, according to the formula below.
      e. Power up the SVM.
      f. Exit MM.

   The memory allocation formula:

Component	Memory allocation rules
Base SVM	• 350 MB
MDM (Master/Slave)	• 470 MB + (500 KB * 8 TB of volume capacity) + (1.44 KB * number of volumes) + (4 KB * number of SDS devices) • Maximum supported volumes: 256 K

Component	Memory allocation rules		
Tie Breaker MDM	• 50 MB		
SDS	• (Base) 536 MB + (RmCache Size) * 1.15 + (Storage capacity in TB) * 53 MB • For SDS high performance profile, we add 195 MB.		
SDC	• 132 KB + 23 MB * (number of MDMs) + 25 KB * (number of SDSs) + 1.5 KB * (number of volumes) + 16 B * (number of volume blocks) + 24 KB * (8 TB of volume capacity) • Volume blocks: 1 GB storage = 8 volume blocks		
RFcache	• 16 * (cache_size/page_size) • Commonly-used sizes:		
	RFcache page size	RFcache memory requirement, if the cache device is 800 GB	RFcache memory requirement, if the cache device is 1.6 TB
	64 K	200 MB	400 MB
	32 K	400 MB	800 MB
	16K	800 MB	1.6 GB
	8 K	1.6 GB	3.2 GB
	4 K	3.2 GB	6.4 GB

Viewing the status of volumes by drv_cfg --rescan command

View the status of volumes configured on a SDC node by running the `drv_cfg --rescan` command.

Before you begin

Ensure that the volumes are created and mapped to the SDC nodes in the ScaleIO system.

In a specific scenario in ScaleIO, volumes may not be visible in the OS, or you may want to view the status of volumes immediately, for any of the following reasons:

- Confirming changes in volume-related configuration
- Viewing volumes immediately
- Confirming inaccuracies viewed in the volume display (for reasons such as temporary communication-related issues between system components)

You may troubleshoot and view the status of volumes immediately by running the `drv_cfg --rescan` command.

Perform the following steps:

Procedure

1. Check if the volume is mapped to an SDC server.

2. Determine if the SDC is installed on the server and connected to an MDM.
3. Ensure that the MDM management IP address is up and running.
4. On an application server, re-scan the volumes, by running the following command:
   - Linux: Run `/opt/emc/scaleio/sdc/bin/drv_cfg --rescan`
   - Windows: Run `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg --rescan`
   - ESXi: Rescan for new devices

   For additional information, see Application server does not see a ScaleIO volume on page 248.

Results

The up-to-date status of volumes is displayed.

Virtual IP feature is not functional

Configuring virtual IP addresses on an MDM cluster via vSphere plug-in may disable the virtual IP feature in the system.

It is recommended that you configure the virtual IP addresses on a physical system only via Installation Manager (IM). The virtual IP configuration using vSphere plug-in may lead to communication issues with SDCs, which may trigger IO errors and cause the virtual IP feature to fail.
However, if you have already used the vSphere plug-in instead, you must perform the following steps to resolve the issue:

Procedure

1. Remove the virtual IP addresses by running the following command:

```
scli --modify_cluster_virtual_ips
```

2. Remove the virtual IP interfaces by running the following command:

```
scli --modify_virtual_ip_interfaces
```

3. Edit the IP addresses in the `drv_cfg` file of each SDC in the system.
   - Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
   - Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
   - ESX: Contact Customer Support for access to this tool.
4. Reconfigure the virtual IP addresses, via IM.

Results

The virtual IP addresses are configured on the MDM cluster in the system, via IM, preventing the failure of the virtual IP feature.

Enabling acceleration in a new node

Enable acceleration in a new SDS node in a ScaleIO system.

Before you begin

Ensure that you have:

- Admin-level username and password for accessing ScaleIO
- The RFcache package (`xcache`) for installation on a new SDS
- This task requires other ScaleIO procedures, that are described in other documents:
  - Installing an SDS in ScaleIO
  - Adding an SDS to the MDM
  - Enabling RFcache, using the Installation Manager (IM)
  - Accessing the SVM

Procedure

1. Install an SDS in the ScaleIO system and add it to the MDM.
2. Enable RFcache on the SDS:

SDS OS	Description
ESXi	Access the SVM and then run the following command: `rpm -i EMC-ScaleIO-xcache-2.0-14000.0.sles11.3.*SVM*.x86_64.rpm`
Linux	Enable RFcache using the Installation Manager.
Windows	Enable RFcache using the Installation Manager.

3. Identify the SSD and HDDs to be accelerated and used as cache devices:

SDS OS	Description
ESXi	Run the esxcli commands or open the vSphere View GUI to check the device properties.
Linux	Run the following commands: a. Run the command `fdisk -l` and check the size of the device in the output. b. Run the command `smartctl` and identify the type of device in the output.
Windows	Run either of the following commands: • `wmic diskdrive list brief /format:list` and check the device properties in the output. • `wmic diskdrive get name,size,model` and check the device name, size, and model number in the output.

4. Log in to the Master MDM.

5. Add an SSD device to the default cache pool:

```
scli --add_sds_rfcache_device (--sds_id <ID> |--sds_name
<NAME> | --sds_ip <IP> [--sds_port <PORT>])
--rfcache_device_path <PATH> [--rfcache_device_name <NAME>]
```

Example:

```
scli --add_sds_rfcache_device --sds_name SDS1 --
rfcache_device_path /dev/sdb --rfcache_device_name
rfcache1_SDS1_sdb
```

6. Repeat the previous steps, from the beginning of the procedure, for each new SDS device to be configured in the RFcache Storage Pool.
7. Enable caching on the Storage Pool in the Protection Domain:

```
scli --set_rfcache_usage (((--protection_domain_id <ID> | --
protection_domain_name <NAME>) --storage_pool_name <NAME>) |
--storage_pool_id <ID>) [--use_rfcache | --dont_use_rfcache]
```

Example:

```
scli --set_rfcache_usage --protection_domain_name PD1 --
storage_pool_name SP1 --use_rfcache
```

8. Verify the status of acceleration per SDS:

```
scli --query_sds (--sds_id <ID> | --sds_name <NAME> |--sds_ip
<IP>)
```

Example:

```
scli --query_sds --sds_ip 192.168.1.6
```

Output, similar to the following, appears, for different command types:

- Command to verify the RFcache status:

```
scli --query_sds (--sds_id <ID> | --sds_name <NAME> |--
sds_ip <IP>)
```

Example:

```
scli --query_sds --sds_ip 9.10.21.141
```

Output:

```
SDS 4eecc17300000002 Name: SDS_10.136.215.203 Version:
2.0.13000
```

```
Protection Domain: 928cd6e100000000, Name:
Default_Protection_Domain
DRL mode: Volatile
Authentication error: None
IP information (total 2 IPs):
         1: 9.10.21.141      Role: All (SDS and SDC)
         2: 9.10.121.141     Role: All (SDS and SDC)
        Port: 7072
RAM Read Cache information:
        128.0 MB (131072 KB) total size
        Cache is enabled
        RAM Read Cache memory allocation state is PENDING
Rfcache enabled
.
.
.

Rfcache device information (total 1 devices):
         1: Name: 50000396CC883494  Path: /dev/sdc  Original-
path: /dev/sdc  Size 762495MB  ID: e177c11700020000
```

- Command to verify Storage Pool status:

```
scli --query_storage_pool (((--protection_domain_id <ID> |
--protection_domain_name <NAME>) --storage_pool_name
<NAME>) | --storage_pool_id <ID>)
```

Example:

```
scli --query_storage_pool --storage_pool_name
Cached_Hdd_Pool_1 --protection_domain_name
Default_Protection_Domain
```

Output:

```
Storage Pool Cached_Hdd_Pool_1 (Id: ada1bb8500000000) has 0
volumes and 23.0 TB (23536 GB) available for volume allocation
        The number of parallel rebuild/rebalance jobs: 2
        Rebuild is enabled and using Limit-Concurrent-IO
policy with the following parameters:
        Number of concurrent IOs per device: 1
        Rebalance is enabled and using Favor-Application-IO
policy with the following parameters:
        Number of concurrent IOs per device: 1, Bandwidth
limit per device: 10240 KB per second
        Background device scanner: Disabled
        Zero padding is enabled
        Spare policy: 36% out of total
        Checksum mode: disabled
        Doesn't use RAM Read Cache
        Uses Flash Read Cache
```

9. Repeat this entire procedure on the rest of the SDS nodes.

Results

The new nodes are configured in the RFcache Storage Pool in ScaleIO.

The VMware plug-in responds slowly

In an environment with a large-scale of volumes, the plug-in response may slow down. To solve this, increase the maximum memory size to 2GB and then restart the vSphere web client service.

- Windows

  In the configuration file, `C:\Program Files\VMware\Infrastructure\vSphereWebClient\server\bin\service\conf\wrapper.conf`, look for a string similar to `= -Xmx` (the line should also start with `wrapper.java.additional`), and change the value to 2048M.

- Linux

  In the configuration file, `usr/lib/vmware-vsphere-client/server/wrapper/conf/wrapper.conf`, look for `wrapper.java.maxmemory=` and change the value to 2048.

  The `wrapper.java.maxmemory` parameter may not exist in vCenter 6.0 `wrapper.conf`, If this is the case, add it manually:

  ```
  wrapper.java.maxmemory=2048
  ```

Solving ScaleIO performance issues

ScaleIO is designed to generate the best performance possible from any given system configuration, by using all possible nodes and distributing the data evenly among them.

When the system performance does not meet your expectations, verify if:

- The relevant volume is allocated to a high-performance Storage Pool.
  For example, is it allocated to a pool consisting of SSDs only? If not, using such a Storage Pool will generate better performance.
- The relevant volume resides in a Storage Pool that consists of different storage drives, with different performances?
  A low performance drive in the pool will slow down all the members (waiting for it to respond). If possible, avoid mixing different types of drives.
- The network in use provides maximal network bandwidth to all the ports in use by ScaleIO.

For a full performance review, use the following resources:

- For all ScaleIO-related products, see *ScaleIO Performance Fine-Tuning Technical Notes*.
- For ScaleIO software only, see *System Analysis Best Practice Guide*.

After following these suggestions, you may contact EMC Support for professional analysis and assistance.

Mismatch in IO counters

Resolving mismatch in the IO counters of ScaleIO Dashboard and customer application.

The IO counter in ScaleIO Dashboard might display a read/write IO value which is larger than the one displayed in the customer application, although the total bandwidth in both the cases is the same.

The mismatch between the IO counters occurs when the customer operating system, which triggers the IO requests to the SDS, splits them into two or more smaller-sized IO requests; the total bandwidth being the same as the bandwidth of the original IO request. Thus the IOs are split at the level of the customer operating system (which causes ScaleIO to see a larger number of IOs), and are integrated as well at the same level (the level that splits them).

Currently, ScaleIO does not support any workaround for such scenarios.

Deploying SVM on a node with a management IP address from a different subnet than the node with the SVM template

Due to a VMware limitation, it is not possible to deploy an SVM on a node with a management IP address that is from a different subnet than the node that hosts the SVM template.

When it is necessary to deploy an SVM on a node with a management IP address that is from a different subnet than the node hosting the SVM template, use one of the following workarounds:

- Add the management IP address from the same SVM subnet to the node that hosts the SVM template, and add the new SVM to an existing system.
- Upload the SVM template to the node that will host the new SVM, and add that node to an existing system. Use the template you uploaded to deploy the SVM.

Speeding up rebuild and rebalance processes

You can set concurrent activity limits to speed up rebuild and rebalance processes.

You can speed up rebuild and rebalance processes. This can be very useful If you are planning maintenance, have a disaster recovery (DR) situation, or must have faster rebuild/rebalance times.

You can use the SCLI or the GUI:

- SCLI - Use the `scli --set_rebuild_rebalance_parallelism` command to increase the amount of concurrent activities.
- GUI - From the **Backend** view, sort by Storage Pool. Right-click the Storage Pool and select **Set I/O Priority**.

Check the service-level agreement (SLA) of your environment/application to ensure that the selected setting does not adversely affect your applications/clients. Test your environment to determine what the optimal "limit" is.

The following example sets concurrent activity to 10:

```
scli --set_rebuild_rebalance_parallelism
--protection_domain_name default --storage_pool_name default
--limit 10 --mdm_ip 10.13.168.138
```

For more information, see the *CLI Reference Guide*.

RFcache is not installed on an upgraded Windows server

When using the Installation Manager (IM) to upgrade ScaleIO from v2.x to a later 2.x version on a Windows server, the LIA removes the old RFcache file (`xcache`) and then reboots the host. After the reboot, the LIA fails to install the new RFcache. This error is displayed in the IM.
To avoid this, before starting the upgrade process from the Installation Manager, perform the following steps:

Note

If the SDC is also installed on the Windows server, you must first remove the SDC, restart the machine, and then remove RFcache.

Procedure

1. Log in to the Master MDM, either directly, or via SSH/RDP.
2. On each relevant Storage Pool, disable RFcache:

```
scli --set_rfcache_usage --protection_domain_name <PD_NAME> --
storage_pool_name <SP_NAME> --dont_use_rfcache
```

3. Remove each RFcache device:

```
scli --remove_sds_rfcache_device --sds_ip <SDS_IP> --
rfcache_device_path <RFCACHE_PATH>
```

4. On all Windows hosts with RFcache, perform the following steps, one host at a time:
   a. From **Control Panel** > **Add/Remove Programs**, uninstall EMC-scaleIO-xcache.
   b. When prompted, reboot the host.
   c. When the host has finished rebooting, clean the registry by running the `Clean_XC_registry.bat` script.

      The script is included in `Complete_Windows_SW_Download.xip`.
5. Upgrade the Gateway.

   From the extracted download file, copy the ScaleIO Gateway MSI to the IM server and run the file: `EMC-ScaleIO-gateway-2.0-14000.X-x64.msi`.
6. After the Gateway is upgraded, when uploading packages to the Gateway, do not upload the RFcache package.
7. Start the upgrade, as described in the relevant section of the upgrade documentation.
8. Once the upgrade is complete, perform an extend operation to add back RFcache.

SSD devices are not recognized in the operating system

SSD devices are not recognized correctly by the ESXi host and are not displayed as SSDs in the vSphere plug-in.

Procedure

1. Select the appropriate method to identify the devices' type and to correspond it to the device ID.

Option	Description

DirectPath is not configured	Perform the following procedures on the ESXi host:
	a. List the devices and their types:
	`cd /opt/lsi/perccli/`
	`./perccli /c0/eall/sall show`
	The devices are displayed with their DG number. The devices' type is displayed in the **Med** column.
	b. Display the DG number and corresponding VD number:
	`./perccli /c0/vall show`
	c. Run the following command:
	`esxcli storage core path list`
	The VD number is listed in the **Target** column. Identify the row with the value in the **naa** column that matches the Target number.

DirectPath is configured	Perform the following procedures on the SVM:
	a. List the devices and their types:
	`/opt/MegaRAID/perccli/perccli64 /c0/eall/sall show`
	The devices are displayed with their DG number. The devices' type is displayed in the **Med** column.
	b. Display the DG number and corresponding VD number:
	`/opt/MegaRAID/perccli/perccli64 /c0/vall show`
	c. Run the following command:
	`ls -l /dev/disk/by-path/`
	Search for "pci-0000:02:00.0-scsi-0:2:*X*:0", where *X* is the VD number from the previous step.

ESX competing thread setting resets on server reboot

The setting for ESXi outstanding IO for competing worlds (threads) on a ScaleIO volume is not persistent, and needs to be applied again after every ESXi host reboot.

You can make the setting persistent by adding a script to the ESXi cron job.

For more information on using a cron job, see https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1033346

Procedure

1. Add the following script on the ESXi host crontab:

```
/etc/rc.local.d/local.sh
#!/bin/sh

line_item=$(esxcli storage core device list|egrep '^((\w+\.)|
( )((Display Name:)|(Devfs Path:)|(Vendor:)|(Model:)|
(Revision:)|(SCSI Level:)|(No of outstanding IOs with
competing worlds:)))(.?)'|sed 's/: $/: Null/g;s/ //g;s/(.)
(worlds\: [0-9]*$)/\1\2\n\n/'|sed -r ':a;N;$!ba;s/\n
{1}

/,/g;s/(\,\,\,)/\n/g;s/\,\,//g;s/ /-/g')

for item in $
{line_item}

; do
new_item=$(echo $
{item}

|sed -r 's/-/ /g;s/(\w+(\s)?\w+)+\:\s//g')
device=$(echo $
{new_item}|cut -d"," -f1);
vendor=$(echo ${new_item}

|cut -d"," -f4);
req_outstanding=$(echo $
{new_item}

|cut -d"," -f8);
if ! [[ "$
{req_outstanding}

" == "256" ]]; then
if [[ "$
{vendor}

" == "EMC" ]]; then
esxcli storage core device set -d $
{device}

-O 256;
fi
fi
done
```

Results

The settings will be persistent.

CHAPTER 15

Frequently Asked Questions

This chapter describes frequently asked questions regarding ScaleIO. Topics include:

- Install the ScaleIO GUI 276
- Associating ScaleIO volumes with physical disks 276
- Port usage and changing default ports 278
- Adding an external SDC to an existing ScaleIO system 279
- Changing the LIA configuration file 281
- Cleaning the ScaleIO VMware environment and performing a clean install 282
- Configuring ScaleIO devices in Linux LVM 283
- Configuring session timeout parameters 284
- Fixing keytool errors 284
- Installing Java on SUSE 12 servers 285
- Mounting ScaleIO 286
- The ScaleIO Gateway web server isn't responding 287
- Upgrading the Gateway when a custom certificate is used 289
- Uploading a new OVA 289
- Using the same data network for different NICs 289
- What to do when the default self-signed certificate expires 289
- Add another IP address subnet to an MDM cluster 290

Install the ScaleIO GUI

You can install the ScaleIO GUI.

Before you begin

- Ensure that the workstation satisfies the requirements described in the "System Requirements" section of the documentation.
- Get the installation file, `ScaleIO GUI for Windows Software Download`, either from the product ISO or the EMC Support Site.

Procedure

1. Install the GUI:

```
EMC-ScaleIO-gui-2.0-14000.X.msi
```

After you finish

To log in to the GUI, see "Logging in to the ScaleIO GUI."

Associating ScaleIO volumes with physical disks

This section describes how to associate volumes with physical disks.

Contact ScaleIO Customer Support for access to the troubleshooting utility.

To get ScaleIO volume information, run the `scli --query_all_volumes` (or `--query_all` or `--query_volume`) command.

Output similar to the following appears:

```
Query-all-volumes returned 10 volumes
Protection Domain 0728185d00000000 Name: pd1
Storage Pool ad99eaab00000000 Name: default
<No volumes defined>

Storage Pool ad99eaac00000001 Name: sp1
Volume ID: fac22a6300000000 Name: vol0 Size: 152.0 GB (155648 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6400000001 Name: vol1 Size: 400.0 GB (409600 MB) Mapped to 1 SDC Thin-provisioned
Volume ID: fac22a6500000002 Name: vol2 Size: 80.0 GB (81920 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6600000003 Name: vol3 Size: 392.0 GB (401408 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6700000004 Name: vol4 Size: 96.0 GB (98304 MB) Mapped to 1 SDC Thin-provisioned
Volume ID: fac22a6800000005 Name: vol5 Size: 112.0 GB (114688 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6900000006 Name: vol6 Size: 96.0 GB (98304 MB) Mapped to 1 SDC Thin-provisioned
Volume ID: fac22a6a00000007 Name: vol7 Size: 176.0 GB (180224 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6b00000008 Name: vol8 Size: 272.0 GB (278528 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6c00000009 Name: vol9 Size: 360.0 GB (368640 MB) Mapped to 1 SDC Thin-provisioned
```

This output shows the Volume ID and name, as well as other volume information.

Volume information - Linux

On the SDC host, run the following command to get the operating system volume information that correlates to the ScaleIO scini device name:

```
ls -l /dev/disk/by-id/ |grep scini
```

Output, similar to the following appears:

```
lrwxrwxrwx 1 root root 12 Aug 25 19:40 emc-vol-62c093a52d14aec7-fac22a6300000000 -> ../../scinia
lrwxrwxrwx 1 root root 12 Aug 25 19:40 emc-vol-62c093a52d14aec7-fac22a6400000001 -> ../../scinic
lrwxrwxrwx 1 root root 12 Aug 25 19:40 emc-vol-62c093a52d14aec7-fac22a6500000002 -> ../../scinib
lrwxrwxrwx 1 root root 12 Aug 25 19:41 emc-vol-62c093a52d14aec7-fac22a6600000003 -> ../../scinie
lrwxrwxrwx 1 root root 12 Aug 25 19:41 emc-vol-62c093a52d14aec7-fac22a6700000004 -> ../../scinid
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6800000005 -> ../../scinif
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6900000006 -> ../../scinig
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6a00000007 -> ../../scinii
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6b00000008 -> ../../scinih
lrwxrwxrwx 1 root root 12 Aug 25 19:43 emc-vol-62c093a52d14aec7-fac22a6c00000009 -> ../../scinij
```

This output shows the scini volume name and the volume ID.

By matching the volume ID in both outputs, you can match the operating system names, sciniX, with the ScaleIO volume name.

For example:

- scinia = fac22a6300000000 = vol0
- scinic = fac22a6400000001 = vol1

Alternatively, run the `sg_inq /dev/sciniX` SCSI query command. The result of this command includes the EMC volume ID at the bottom of the output, as illustrated in the following figure:

```
Vendor identification: EMC
Product identification: ScaleIO
Product revision level: 1.3
Unit serial number: EMC-62c093a52d14aec7-fac22a6300000000
```

Note

The sg3_utils must be installed on the Linux host in order to run this command.

Volume information - Windows

The `sg_inq.exe` file was added to the MSI installation and can be found at `C:\Program Files\EMC\ScaleIO\SDC\diag\`.

Procedure

1. Run the `sg_inq HardiskX` SCSI query command.

   The result of this command includes the EMC volume ID at the bottom of the output.

2. On the MDM, get the ScaleIO volume information:

   ```
   C:\Program Files\emc\scaleio\sdc\bin\drv_cfg --query_vol
   ```

   Output similar to the following is displayed:

   ```
   Retrieved 5 volume(s)
   VOL-ID 6acb988100000000 MDM-ID 0b246c9a755ca3dd
   VOL-ID 6acb988200000001 MDM-ID 0b246c9a755ca3dd
   VOL-ID 6acb988300000002 MDM-ID 0b246c9a755ca3dd
   VOL-ID 6acb988400000003 MDM-ID 0b246c9a755ca3dd
   VOL-ID 6acb988500000004 MDM-ID 0b246c9a755ca3dd
   ```

3. From the Windows command prompt, run this command:

```
wmic diskdrive get deviceid,serialnumber | findstr "EMC"
```

Output similar to the following is displayed:
`\\.\PHYSICALDRIVE13 EMC-0b246c9a755ca3dd-6acb988500000004`

The first part of the output is the disk name. In our example:
`PHYSICALDRIVE13`

The second part is the disk serial number. The last set of the second part (after the dash) is the ScaleIO volume ID. In our example: `6acb988500000004`

After you finish

You can also get the volume ID from the ScaleIO GUI by displaying the **Identity** pane of the volume's properties sheet from **Frontend** > **Volumes**

Port usage and changing default ports

The following table lists the TCP ports that are used by ScaleIO. Prior to installing or upgrading a system, ensure that these ports are not in use by other processes.

If they are in use, either free them or change them to another available port.

Table 31 Default ports

Port used by	Port #	Protocol	File to change	Field to modify (or to add, if it does not exist)	Notes
MDM listener	6611	Protobuf over TCP	**Note** Cannot be modified, and must be available		
MDM Cluster member	9011	Protobuf over TCP	`/opt/emc/scaleio/mdm/cfg/conf.txt`	`actor_cluster_port=`*<NEW_PORT>*	
SDS listener	7072	Proprietary protocol over TCP	`/opt/emc/scaleio/sds/cfg/conf.txt`	`tgt_port=`*<NEW_PORT>*	SDCs connect through this port for data communication and to the MDM for meta-data communication. When multiple SDSs are installed on the same physical server, use ports 7072+*x*, where *x* is the index of the SDS (for example, 70721, 70722).
LIA listener	9099	Protobuf over TCP	`/opt/emc/scaleio/lia/cfg/conf.txt`	`lia_port=`*<NEW_PORT>*	The Installation Manager connects to the LIA to perform installation and maintenance-related operations.
Gateway-Installation Manager/REST (not secure)	80 (or 8080, together with 8443)	REST over HTTPS	`<gateway installation directory>/conf/catalina.properties`	`http.port=`80 (or 8080)	After changing the port, you must restart the service/daemon: • Linux: Run `service scaleio-gateway restart` • Windows: Restart the EMC ScaleIO Gateway service

Table 31 Default ports (continued)

Port used by	Port #	Protocol	File to change	Field to modify (or to add, if it does not exist)	Notes
Gateway-Installation Manager/REST (secure)	443 (or 8443, together with 8080)	REST over HTTPS	`<gateway installation directory>/conf/catalina.properties`	`ssl.port`=443 (or 8443)	
SNMP	162	SNMP v2 over UDP			SNMP traps for system alerts are sent to a trap receiver via this port. The ScaleIO gateway sends messages to: snmp.traps_receiver_ip on the port snmp.port
SDBG for MDM (Manager)	25620				Used by ScaleIO internal debugging tools to extract live information from the system for debugging purposes. When multiple SDSs are installed on the same physical server, use ports 2564+*x*, where *x* is the index of the SDS (for example, 25641, 25642).
SDBG for MDM (Tie Breaker)	25600				
SDBG for SDS	25640				

Adding an external SDC to an existing ScaleIO system

During manual installation, you can install the SDC according to the OS-specific instructions in the following section, and it will be connected to the existing ScaleIO system:

Installing SDC on an ESX server and connecting it to ScaleIO

Install the SDC with the appropriate parameters to connect it to an existing ScaleIO system.

Before you begin

Ensure that you have:

- The virtual IP address or MDM IP address of the existing system
- Login credentials for the SDC
- The appropriate installation packages for the SDC
- Access to the drv_cfg tool. Contact EMC support for access to this tool on ESX.

The following procedure describes installing an external SDC on an ESX server using the esxcli. Alternatively, you can install the external SDC using the vSphere plug-in. For more information, see "Installing the SDC on ESX hosts" in the *ScaleIO Deployment Guide*.

Procedure

1. On the ESX on which you are installing the SDC, set the acceptance level:

```
esxcli --server=<SERVER_NAME> software acceptance set --
level=PartnerSupported
```

   where *<SERVER_NAME>* is the ESX on which you are installing the SDC.

2. Install the SDC:

```
esxcli software vib update –d “Full Path”
```

3. Set the IP address of the MDM:

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=<XXXXXX> IoctlMdmIPStr=<LIST_VIP_MDM_IPS>"
```

   where

   - *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM
   - *<XXXXXX>* is the version

Results

The SDC is installed on the ESX server and is connected to the ScaleIO system.

Installing SDC on a Linux server and connecting it to ScaleIO

Install the SDC with the appropriate parameters to connect it to an existing ScaleIO system.

Before you begin

Ensure that you have:

- The virtual IP address or MDM IP address of the existing system
- Login credentials for the SDC
- The appropriate installation packages for the SDC

The following procedure describes manually installing an external SDC on a Linux server. Alternatively, you can install the external SDC using the Installation Manager. For more information, see "Extending an existing ScaleIO system" in the *ScaleIO Deployment Guide*.

Procedure

1. Install the SDC:
   - RHEL/CentOS /Oracle Linux

```
MDM_IP=<LIST_VIP_MDM_IPS> rpm -i <SDC_PATH>.rpm
```

   - Ubuntu

```
MDM_IP=<LIST_VIP_MDM_IPS> dpkg -i <SDC_PATH>.deb
```

- CoreOS

```
MDM_IP=<LIST_VIP_MDM_IPS> ./<LIST_VIP_MDM_IPS>.bsx
```

where

- *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM
- *<SDC_PATH>* is the path where the SDC installation package is located

Results

The SDC is installed on the Linux server and is connected to the ScaleIO system.

Installing SDC on a Windows server and connecting it to ScaleIO

Install the SDC with the appropriate parameters to connect it to an existing ScaleIO system.

Before you begin

Ensure that you have:

- The virtual IP address or MDM IP address of the existing system
- Login credentials for the SDC
- The appropriate installation packages for the SDC

The following procedure describes manually installing an external SDC on a Windows server. Alternatively, you can install the external SDC using the Installation Manager. For more information, see "Extending an existing ScaleIO system" in the *ScaleIO Deployment Guide*.

Procedure

1. On the Windows server on which you are installing the SDC, run:

```
msiexec /i <SDC_PATH>.msi MDM_IP=<LIST_VIP_MDM_IPS>
```

where

- *<SDC_PATH>* is the path where the SDC installation package is located
- *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM

Results

The SDC is installed on the Windows server and is connected to the ScaleIO system.

Changing the LIA configuration file

You can change the default behavior of the LIA by editing the configuration file:

- Windows: `C:\Program Files\emc\scaleio\LIA\cfg\conf.txt`
- Linux: `/opt/emc/scaleio/lia/cfg/conf.txt`

  The following are some values relevant to LIA behavior:

```
lia_token=5
lia_enable_install=1
```

```
lia_enable_uninstall=1
lia_enable_configure_fetch_logs=1
```

For example, to restrict which Gateway IP addresses can access the LIA, add those IP addresses to this line in the `conf.txt` file:

```
lia_trusted_ips=<IP_ADDRESS_1>,<IP_ADDRESS_2>
```

To set this during LIA installation, set the TRUSTED_IPS environment variable. For example:

```
TRUSTED_IPS=1.2.3.4,5.6.7.8 rpm -i lia.rpm
```

Cleaning the ScaleIO VMware environment and performing a clean install

This topic explains how to clean the ScaleIO VMware environment and perform a clean install while using previously defined networks.

Before you begin

Before you begin, unmap and delete any ScaleIO volumes in your system.

If necessary, unregister your ScaleIO system from within the plugin and delete all the ScaleIO SVMs.

Procedure

1. Set to **Run as administrator**, close the existing PowerCLI sessions and open a new one.
2. Using the PS1 script, unregister the plugin.
3. Stop the vSphere web client service:

   VC Linux: `service vsphere-client stop`
4. Delete the contents of the plug-in folder.

   The vSphere web client (Virgo) plug-in folders are located at:

vCenter	Operating system	Path to file
5.x	Windows	C:\ProgramData\VMware\vSphere Web Client\vc-packages\vsphere-client-serenity
	Linux	/var/lib/vmware/vsphere-client/vc-packages/vsphere-client-serenity/
6.x	Windows	C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client\vc-packages\vsphere-client-serenity
	Linux	/etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity

5. Delete the `scaleio` folder or its contents.

   The `scaleio` folders are located at:

vCenter	Operating system	Path to file
5.x	Windows	C:\Windows\System32\config\systemprofile\AppData\Roaming\VMware\scaleio
	Linux	/opt/.vmware/scaleio
6.x	Windows	C:\Users\vspherewebclientsvc\AppData\Roaming\VMware\scaleio
	Linux	/etc/vmware/vsphere-client/vc-packages/scaleio

6. Clean the Virgo logs folder.

   The Virgo log folders are located at:

vCenter	Operating system	Path to file
5.x	Windows	C:\ProgramData\VMware\vSphere Web Client\serviceability\logs
	Linux	/var/log/vmware/vsphere-client/
6.x	Windows	C:\ProgramData\VMWare\vCenterServer\logs\vsphere-client\logs
	Linux	/var/log/vmware/vsphere-client/logs

7. Start the vSphere web client service:

   VC Linux: `service vsphere-client start`

8. Clear your web browser's cache and cookies, or else open a different web browser.
9. Using the PS1 script, register the plugin via PowerCLI.

   Note

   Do not press ENTER at this point.

10. After you have logged in to the vSphere web client to complete the registration and you see the ScaleIO icon, press ENTER in the PowerCLI session.

    This stops the embedded Tomcat server.

11. If necessary, remove the SDC module parameters and VIB from the ESXs:

    a. Connect via SSH to each ESX.

    b. Run:

```
~ # esxcli system module parameters set -m scini -p ""
~ # esxcli software vib remove -n scaleio-sdc-esx5.5 / 6.0
```

    c. Reboot each ESX.

Configuring ScaleIO devices in Linux LVM

To configure ScaleIO devices, perform the following:

Procedure

1. Edit the `/etc/lvm/lvm.conf` file by adding the following line:

```
types = [ "scini", 16 ]
```

2. If only ScaleIO scini devices are to be used, you can add the following filter:

```
filter = [ "a|/dev/scini*|", "r/.*/" ]
```

3. Once configured, the `lvmdiskscan` command should yield results similar to the following:

```
/dev/scinia  [      96.00 GiB] LVM physical volume
/dev/scinib  [     320.00 GiB] LVM physical volume
/dev/scinic1 [      56.00 GiB]
/dev/scinid  [      32.00 GiB]
1 disk
1 partition
2 LVM physical volume whole disks
0 LVM physical volumes
```

4. Continue with normal LVM steps.

Configuring session timeout parameters

When a user is authenticated by the system, all commands are performed with the user's respective role until a logout is performed, or until the session expires by reaching one of the following timeouts:

- Maximum session length (default: 8 hours)
- Session idle time (default: 10 minutes)

You can modify these parameters, by editing the MDM `conf.txt` file:

- Linux: `/opt/emc/scaleio/mdm/cfg/conf.txt`
- Windows: `C:\Program Files\emc\scaleio\mdm\cfg\conf.txt`

1. To configure maximum session length, edit the value of the `user_session_hard_timeout_secs` parameter. The minimum is 10 seconds, maximum 10 years, and default 8 hours.
2. To configure session idle time, edit the value of the `user_session_timeout_secs` parameter. The minimum is 10 seconds, maximum 3 months, default 10 minutes.
3. After changing the parameters, restart the MDM service (delete and create service) for the changes to take effect.
4. To ensure persistence after MDM restart, make these changes on every MDM.

Fixing keytool errors

Error during rpm installation command

Error message:

```
No keytool path was found. Please pass SIO_GW_KEYTOOL as an argument
to the rpm installation command.
```

If a message similar to this is displayed after executing the rpm command to install the Gateway, add the location of the `/bin/keytool` file on your server to the command.

Example:

```
SIO_GW_KEYTOOL=/usr/lib/jvm/java-1.6.0-openjdk-1.6.0.0.x86_64/jre
rpm -U <gateway_installation_file_name>.rpm
```

Error during rpm upgrade command

Error message:

```
No keytool path was found. Set the environment variable SIO_GW_KEYTOOL
```

If a message similar to this is displayed after executing the rpm command to upgrade the Gateway, add the location of the `/bin/keytool` file on your server to the command.

Example:

```
SIO_GW_KEYTOOL=/usr/java/default/bin/ rpm -U /tmp/EMC-ScaleIO-
gateway-1.32-363.0.x86_64.rpm
```

Installing Java on SUSE 12 servers

Installation of Java is different in SLES-based distributions because SLES uses update-alternatives commands. For SUSE, we use a TGZ file in place of RPM.

To install Java on SUSE 12 servers:

Procedure

1. Untar the TGZ (for example, `jre-8u60-linux-x64.tar.gz`) to `/usr/java`.

   This creates a directory of `/usr/java/jre1.8.0_60/`.

2. Apply the std update-alternatives procedure:

   ```
   /usr/sbin/update-alternatives --install "/usr/bin/java"
   "java" "/usr/java/jre1.8.0_60/bin/java" 40
   /usr/sbin/update-alternatives --config java
   /usr/sbin/update-alternatives --install "/usr/bin/keytool”
   "keytool" "/usr/java/jre1.8.0_60/bin/keytool" 40
   /usr/sbin/update-alternatives --config keytool
   ```

Mounting ScaleIO

The exposed ScaleIO volumes are connected to the servers via the network. To configure mounting options of ScaleIO devices, follow the instructions for your operating system.

Use persistent device names, described in full in Associating ScaleIO volumes with physical disks on page 276.

To mount ScaleIO:

Procedure

1. Determine the `/dev/disk/by-id` correlation to `/dev/sciniX`:

```
ls -l /dev/disk/by-id/ |grep scini
```

Output similar to the following appears:

```
lrwxrwxrwx 1 root root 12 Mar 2 05:35 emc-
vol-7ec27ef55b8f2108-85a0f0330000000a -> ../../scinia
lrwxrwxrwx 1 root root 12 Mar 2 05:35 emc-
vol-7ec27ef55b8f2108-85a0f03200000009 -> ../../scinib
lrwxrwxrwx 1 root root 12 Mar 2 05:35 emc-
vol-7ec27ef55b8f2108-85a0f02c00000003 -> ../../scinic
```

2. Run the mount command:

```
mount /dev/disk/by-id/<EMC-vol-id>
```

Example:

```
mount /dev/disk/by-id/emc-
vol-7ec27ef55b8f2108-85a0f0330000000a /mnt_scinia
```

3. To make the mount command persistent, edit the `/etc/fstab` file according to the instructions for your operating system:
   - RHEL 6.x:
     a. In `/etc/fstab`, use a text editor to add the ScaleIO mount lines:

```
/dev/disk/by-id/emc-
vol-7ec27ef55b8f2108-85a0f0330000000a /mnt_scinia ext4
defaults 0 0
```

     b. In `/etc/rc.local`, use a text editor to add the mount commands:

```
mount /mnt_scinia
```

   - RHEL 7.x:
     In `/etc/fstab`, use a text editor to add `_netdev` to the ScaleIO mount lines.

Example:

```
/dev/disk/by-id/emc-vol-7ec27ef55b8f2108-85a0f0330000000a /
mnt_scinia ext4 defaults,_netdev 0 0
```

Ensure that you comply with the `netdev` and syntax rules for your file system, as described in the `man` page.

- SLES:
  In `/etc/fstab`, use a text editor to add `nofail` to the ScaleIO Ready Node mount lines.

  Example:

```
/dev/disk/by-id/emc-vol-7ec27ef55b8f2108-85a0f0330000000a /
mnt_scinia ext3 nofail 0 0
```

  Ensure that you comply with the `nofail` and syntax rules for your file system, as described in the `man` page.

The ScaleIO Gateway web server isn't responding

The ScaleIO Gateway (REST service, Installation Manager) may be disabled:

The ScaleIO Gateway seems to be locked or disabled, and returns the HTTP status code 401 or 403.

Solution

- Ensure that the Gateway is enabled, as described in the documentation.
- In the `gatewayUser.properties` file, ensure that the `gateway-admin.password` property has a non-blank password. If the password is blank, the gateway has been locked.

  The following table shows the location of the gatewayUser.properties file:

Gateway installed on	Location of gatewayUser.properties file
Windows, 64-bit	C:\Program Files\EMC\ScaleIO\Gateway\webapps\ROOT\WEB-INF\classes\
Linux	/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes

  To reset the Scaleio-Gateway password, perform the following steps:

Procedure

1. Use SioGWTool to reset the password by typing the following command:

```
SioGWTool --reset_password --password <new_scaleio-
gateway_password> --config_file
<path_to_file_gatewayUser.properties>
```

Note

The path to SioGWTool is:

Linux: `/opt/emc/scaleio/gateway/bin/SioGWTool.sh`

Windows: `C:\Program Files\EMC\ScaleIO\Gateway\bin\SioGWTool.bat`

2. Restart the scaleio-gateway service

The ScaleIO Gateway web server isn't responsive and the following error appears in the catalina log file:

- Windows:

  `C:\Program Files\EMC\ScaleIO\Gateway\logs\catalina.<date>.log`

- Linux:

  `/opt/emc/scaleio/gateway/logs`

```
2014-06-21 22:50:57,113 [main] ERROR
o.a.coyote.http11.Http11NioProtocol - Failed to initialize end
point associated with ProtocolHandler ["http-nio-443"]
java.net.BindException: Address already in use: bind
```

Solution

Perform one of the following:

Procedure

1. Find the service/daemon that is currently occupying that port and stop it:
   - Windows

     Run: netstat -anb
   - Linux

     Run: `netstat -alp`

   On Windows, one of the common applications that occupies this port is the VMware workstation, which uses this port for the shared VM feature. You can configure VMware workstation to use a different port via the Settings dialog, or you can disable the shared VM feature.

   Once the port is free, restart the scaleio-gateway service:

   - Windows

     Restart the EMC ScaleIO Gateway service.
   - Linux

     Type the command `service scaleio-gateway restart`
2. Change the ScaleIO Gateway web server to run on a different port, as described in "Changing default ports".

After doing so, restart the ScaleIO Gateway service/daemon, as described above. Access the Gateway with the new port. For example: https://<host>:<port>

Upgrading the Gateway when a custom certificate is used

If a custom security certificate is used on the ScaleIO Gateway (Windows and Linux environments), you must save a copy of the certificate (`*.keystore` file) and the `catalina.properties` file before you upgrade the gateway. After the upgrade is complete, you must copy these files back to their original location.

The default file locations, per operating system, are:

Linux:

`/opt/emc/scaleio/gateway/conf/catalina.properties`

`/opt/emc/scaleio/gateway/conf/certificates/.keystore`

Windows (64 bit):

`C:\Program Files\EMC\ScaleIO\Gateway\conf\catalina.properties`

`C:\Program Files\EMC\ScaleIO\Gateway\conf\certificates\.keystore`

Uploading a new OVA

If you have already used the OVA to create a template, you cannot create another template with the same name in the same datacenter.

Either remove the original template first, or use the `ScaleIOPluginSetup-2.0-14000.X.ps1` script, option #3, to assign a different name to the new template.

You can also upload the OVA manually using the VMware OVA upload tools. Configure the networks manually, after deployment or during the wizard menus. For more information, see the VMware user guides.

Using the same data network for different NICs

This configuration is supported, but it could reduce efficiency of outgoing communication and deny you the benefits of high availability of the multiple networks.

What to do when the default self-signed certificate expires

If the default self-signed security certificate is used on the ScaleIO Gateway, it expires after approximately one year. When you upgrade the gateway, the self-signed certificate is automatically replaced with a new one. If your self-signed security certificate expires, you can create a new one using the Java keytool utility.

Add another IP address subnet to an MDM cluster

Add an IP network to an existing MDM cluster.

Before you begin

This topic explains how to add another IP address subnet for use by the MDM cluster. This procedure addresses scenarios where the MDM cluster uses a single network, or when an existing network needs to be replaced by a different one.

Note

This procedure describes an example for a for 3-node cluster, however, the procedure for a 5-node cluster is similar.

Procedure

1. Query the system to get the current cluster state/health:

```
scli --query_cluster
```

   Cluster status is returned, where you can identify the Master, the Slave, and the Tie Breaker.

2. Switch to single cluster mode:

```
scli --switch_cluster_mode --cluster_mode 1_node --
remove_slave_mdm_id <mdm_slave_id> --remove_tb_id <tb_id>
```

3. Remove the standby MDM:

```
scli --remove_standby_mdm --remove_mdm_id <mdm_slave_id>
```

4. Remove the Tie Breaker:

```
scli --remove_standby_mdm --remove_mdm_id <tb_id>
```

5. Add the MDM as standby with its IP addresses (including the additional IP addresses):

```
scli --add_standby_mdm --new_mdm_ip ip_1<,ip_2,...> --
mdm_role manager --new_mdm_management_ip ip_1<,ip_2,...> --
allow_asymmetric_ips --force_clean
```

   For example:

```
scli --add_standby_mdm --new_mdm_ip 10.89.9.6,10.89.11.6 --
mdm_role manager --new_mdm_management_ip 10.89.9.6,10.89.11.6
--allow_asymmetric_ips --force_clean
```

6. Add the Tie Breaker as standby with its IP addresses (including the additional IP addresses):

```
scli --add_standby_mdm --new_mdm_ip ip_1<,ip_2,...> --
mdm_role tb --new_mdm_management_ip ip_1<,ip_2,...> --
allow_asymmetric_ips --force_clean
```

7. Switch cluster operation back to a 3-node cluster:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_id <slave_id> --add_tb_id <tb_id>
```

For example:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_id 0x4520631c7262bbf1 --add_tb_id
0x3cde0ef516f61162
```

8. Query the system to get the current cluster state/health.

```
scli --query_cluster
```

Cluster status is returned, where you can check that the cluster is configured and operating as expected.

9. Switch MDM ownership to verify cluster functionality:

```
scli --switch_mdm_ownership --new_master_mdm_id
<new_master_mdm_id>
```

For example:

```
scli --switch_mdm_ownership --new_master_mdm_id
0x4520631c7262bbf1
```

10. Query the system to get the current cluster state/health.

```
scli --query_cluster
```

Cluster status is returned, where you can check that the cluster is operating as expected.

11. Add IP addresses for the Master MDM (presently Slave MDM) by following steps 2, 3, 5, 7, and 8.
12. Optional: Switch MDM ownership back to the original MDM:

```
scli --switch_mdm_ownership --new_master_mdm_id MDM_ID
```

CHAPTER 16

System events

This chapter describes ScaleIO system events and alerts. Topics include:

- System events overview 294
- Event format 294
- Viewing events 295
- Event list 298

System events overview

This appendix describes event messages that can be generated by the ScaleIO system.

An event message is generated as a response to changes that have occurred while the system is running. Event messages to notify you of the changes in case your intervention is needed. Each event message is associated with a severity level. The severity indicates the risk (if any) to the system, in relation to the changes that generated the event message. The severity levels are as follows:

- Info
  Informs you of events that one should be aware of, but that do not put the system at risk (No Urgency)

  Example: `CLI_COMMAND_RECEIVED`
- Warning
  Indicates a failure that may result from an acceptable condition (e.g. user error), but can also indicate a possible failure.

  Example : `SDC_DISCONNECTED`.

  If the disconnection is planned, or self-recovered, then all is OK. Otherwise, this might require user intervention.
- Error
  An error alarm was raised by the system. This error requires your attention. The system is stable, but could be degraded.

  Example: `MDM_DATA_DEGRADED`

  The system is operational but some data is not protected. The system is recovering, but hardware replacement might be required.
- Critical
  A major error alarm was raised by system. The system requires immediate attention.

  Example: `MDM_DATA_FAILED`

  Some data is unavailable.

Event entries are documented as follows:

- Name
  The name associated with the event
- Message
  The message that will appear
- Severity
  The severity level
- Description
  A description of the reasons that triggered the event notification
- Action
  Possible actions that can resolve the reported event (if relevant)

Event format

All event messages received are in a parsable structured format, containing the following fields:

ID: a sequential number attached to all events

Date: the local time set in the server

Format: YYYY-MM-DD hh:mm:ss.ssssss

Name: the unique name of the event

Severity: one of the predefined severity levels

Message: message describing the event. Some event notification message verbosity may be expanded by using the full switch (see "Viewing events locally").

The following is an example of a possible event notification:

```
139  2013-07-22 17:21:11.694571 CLI_COMMAND_RECEIVED
INFO     Command MAP_VOL_TO_SCSI_INITIATOR Received
```

This event has the option of extended verbosity. When requested, the event notification will be displayed as follows:

```
139  2013-07-22 17:21:11.694571 CLI_COMMAND_RECEIVED
INFO     Command MAP_VOL_TO_SCSI_INITIATOR Received Vol
Name: snap_raw; SCSI Initiator Name: ini-21
bAllocateLunNum: 1 lunNum: 0
```

The following is a breakdown of the event according to the fields in the event record (as described above):

Parameter	Description
ID	139
Date	2013-07-22 17:21:11.694571
Name	CLI_COMMAND_RECEIVED
Severity	INFO
Message	Command MAP_VOL_TO_SCSI_INITIATOR Received
Extended	Command MAP_VOL_TO_SCSI_INITIATOR Received Vol Name: snap_raw; iSCSI Initiator Name: ini-21 bAllocateLunNum: 1 lunNum: 0

Viewing events

You can view events in the following ways:

- on a local server
- via Syslog
- via email

To configure events via email, see the *EMC Secure Remote Services Installation and Operations Guide*.

Viewing events locally

Events can be viewed by running the following command, and by using switches to filter the data.

Command:

showevents.py

Syntax

```
/opt/emc/scaleio/mdm/bin/showevents.py [Options]
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

Displays events, which can be filtered by optional switches.

Parameters

Parameter	Description
Options:	
--min_severity <SEVERITY>	Displays events with at least the specified minimum severity
--severity <SEVERITY>	Displays events with the specified severity
--from_id <ID>	Displays all events starting from the given ID
--to_id <ID>	Displays all events ending at the given ID
--from_date <ID>	Displays all events starting from the given date
--to_date <ID>	Displays all events ending at the given date
--grep <TEXT>	Displays events containing the specified text
--full	Extends message verbosity

Example

```
/opt/emc/scaleio/mdm/bin/showevents.py --severity ERROR
```

Viewing events in Syslog

The MDM syslog service can send ScaleIO events, via TCP/IP, to RFC 6587-compliant remote (or local) Syslog servers. Messages are sent with facility `local0`, by default.

Once the syslog service is started, all events will be sent until the service is stopped.

This section describes how to use the CLI to start, stop, and configure the facility field of the syslog events.

Start posting events to remote syslog servers

Command:

start_remote_syslog

Syntax

```
scli --start_remote_syslog --remote_syslog_server_ip
<IP> [--remote_syslog_server_port <PORT>]
[--syslog_facility <FACILITY>] [--attach_event_code]
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

Starts posting events to one or more remote syslog servers.

Parameters

Parameter	Description
--remote_syslog_server_ip <IP>	A comma-separated list of syslog server IP addresses (maximum of two servers)
--remote_syslog_server_port <PORT>	The syslog server port (default 1468)
--syslog_facility <FACILITY>	Controls the facility field of the event (legal values 0—23; default 16)
--attach_event_code	Add the posted event code to the event message (disabled, by default)

Example

```
scli --start_remote_syslog --remote_syslog_server_ip 192.168.1.10 --
remote_syslog_server_port 1500
```

Stop posting events to remote syslog servers

Command:

stop_remote_syslog

Syntax

```
scli --stop_remote_syslog --remote_syslog_server_ip
```

<IP>

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes

Stops posting events to remote syslog servers.

Parameters

Parameter	Description
--remote_syslog_server_ip <IP>	A comma-separated list of syslog servers IP addresses

Example

```
scli --stop_remote_syslog --remote_syslog_server_ip
192.168.1.10,192.168.1.20
```

Configure the syslog events facility field

Command:

`set_syslog_facility`

Syntax

```
scli --set_syslog_facility --remote_syslog_server_ip <IP> --
syslog_facility <FACILITY>
```

Actual command syntax is operating-system dependent. For more information, see "CLI basics".

Description/Notes
Configures the facility field of the syslog events.

Parameters

Parameter	Description
--remote_syslog_server_ip <IP>	A comma-separated list of syslog server IP addresses
--syslog_facility <FACILITY>	Controls the facility field of the event (legal values 0—23; default 16)

Example

```
scli --set_syslog_facility --remote_syslog_server_ip
192.168.1.10,192.168.1.20 --syslog_facility 20
```

Event list

This section lists all ScaleIO events, grouped by the following categories:

- Authentication
- CLI commands
- License and installation
- MDM
- SDC
- SDS
- Rebuild

Authentication

Authentication Failed

Parameter	Description
Name	AUTHENTICATION_FAILED
Message	Authentication failed for user U
Severity	Warning
Description	User entered the wrong password
Action	If you see this event multiple times, someone may be trying to gain unauthorized access to the system.

CLI commands

CLI Command Received

Parameter	Description
Name	CLI_COMMAND_RECEIVED
Message	Command X Received
Severity	Info
Description	CLI command X was entered by a user
Action	None

CLI Command Succeeded

Parameter	Description
Name	CLI_COMMAND_ SUCCEEDED
Message	Command X ended successfully
Severity	Info
Description	CLI command X was executed successfully
Action	None

CLI Command Failed

Parameter	Description
Name	CLI_COMMAND_FAILED
Message	Command X failed with error E
Severity	Warning

Parameter	Description
Description	The CLI command X entered by user failed with error E
Action	Look up error E and address the issue accordingly

Snapshot volumes could not be found, by ID

Parameter	Description
Name	SNAPSHOT_VOLUMES_FAILED_BY_ID
Message	Could not snapshot volumes, because a volume ID was not found. ID: "MOS_OBJID__FORMAT"."
Severity	Error
Description	This message is posted if a snapshot_volume command contains an invalid volume ID (out of many). The CLI will only get an error code, but in the event, you can see which volume ID is invalid.
Action	Verify the parameters entered for the snapshot_volume command

Snapshot volumes could not be found, by name

Parameter	Description
Name	SNAPSHOT_VOLUMES_FAILED_BY_NAME
Message	Could not snapshot volumes, because a volume was not found: Name: %s.
Severity	Error
Description	This message is posted if a snapshot_volume command contains an invalid volume name (out of many). The CLI will only get an error code, but in the event, you can see which volume name is invalid.
Action	Verify the parameters entered for the snapshot_volume command

License and installation

License Expiration Warning

Parameter	Description
Name	LICENSE_EXPIRATION_WARNING
Message	License will expire in X days
Severity	Warning
Description	System license will expire in 30 days or less
Action	Contact EMC Support for license renewal, and then reinstall.

License Expiration Error

Parameter	Description
Name	LICENSE_EXPIRATION_ERROR
Message	License will expire in X days
Severity	Error
Description	System license will expire in 7 days or less
Action	Contact EMC Support for license renewal. If you have already renewed your license, install it.

License Expiration Critical

Parameter	Description
Name	LICENSE_EXPIRATION_CRITICAL
Message	License will expire in X days
Severity	Critical
Description	System license will expire in 2 days or less
Action	Contact EMC Support for license renewal. If you have already renewed your license, install it.

License Expired

Parameter	Description
Name	LICENSE_EXPIRED
Message	License has expired
Severity	Critical

Parameter	Description
Description	The system's license has expired
Action	To resume operational mode, contact EMC Support for license renewal. If you have already renewed your license, install it.

Upgrade has started

Parameter	Description
Name	UPGRADE_STARTED
Message	Upgrade to version %s has started.
Severity	Info
Description	An upgrade procedure has been initiated
Action	Not needed

Upgrade has finished

Parameter	Description
Name	UPGRADE_FINISHED
Message	Upgrade completed successfully.
Severity	Info
Description	An upgrade procedure completed successfully
Action	Not needed

Upgrade has failed

Parameter	Description
Name	UPGRADE_FAILED
Message	Upgrade was not successful. Reason: %s
Severity	Error
Description	An upgrade procedure was not able to complete
Action	Fix the error and retry the upgrade

MDM

MDM Started

Parameter	Description
Name	MDM_STARTED
Message	MDM Process started
Severity	Info
Description	MDM process has started running
Action	None

MDM Data Degraded

Parameter	Description
Name	MDM_DATA_DEGRADED
Message	Some of the Storage Pool data is now in Degraded state
Severity	Error
Description	Some of the Storage Pool data is in Degraded state. This data is not protected against another failure.
Action	The system is rebuilding the Degraded data to return to Normal (protected) state. Check if any hardware is malfunctioning and requires replacement.

MDM Data Failed

Parameter	Description
Name	MDM_DATA_FAILED
Message	Some Storage Pool data is now unavailable
Severity	Critical
Description	Multiple failures have occurred. Some Storage Pool data is now unavailable. This data cannot be accessed.
Action	Locate and fix the failed hardware. If the problem is not resolved, contact EMC Support.

MDM Data Normal

Parameter	Description
Name	MDM_DATA_NORMAL
Message	All of the Storage Pool data has returned to Normal state
Severity	Info
Description	All Storage Pool data previously in Degraded or Failed state has returned back to Normal state. User data is fully accessible and protected.
Action	None

SDC

New SDC Connected

Parameter	Description
Name	NEW_SDC_CONNECTED
Message	New SDC (IP: X; ID: Y; GUID: Z) connected
Severity	Warning
Description	A new SDC (IP: X; ID: Y; GUID: Z) has connected to the MDM
Action	A new SDC has just connected to the MDM. Validate that this is a valid SDC.

SDC Connected

Parameter	Description
Name	SDC_CONNECTED
Message	SDC (IP: X; ID: Y; GUID: Z) reconnected
Severity	Info
Description	An existing SDC (IP: X; ID: Y; GUID: Z) has reconnected to the MDM
Action	None

SDC Disconnected

Parameter	Description
Name	SDC_DISCONNECTED
Message	SDC (IP: X; ID: Y; GUID: Z) disconnected
Severity	Warning
Description	SDC (IP: X; ID: Y; GUID: Z) has disconnected from the MDM.
Action	Make sure it is expected otherwise this might be a hardware malfunction

SDS

SDS Disconnected

Parameter	Description
Name	SDS_DISCONNECTED
Message	SDS X (IP:Y; ID: Z) disconnected
Severity	Error
Description	SDS X (IP: Y; ID Z) has disconnected from the MDM
Action	Make sure that this is an expected event, because otherwise this might be caused by a hardware malfunction.

SDS Reconnected

Parameter	Description
Name	SDS_RECONNECTED
Message	SDS X (IP:Y; ID: Z) reconnected
Severity	Info
Description	SDS X (IP: Y; ID Z) has reconnected to the MDM. If this event appears multiple times subsequently for the same SDS (and not directly after SDS_DISCONNECTED), it can indicate a bad network connection.
Action:	Check network connections.

SDS Remove Done

Parameter	Description
Name	SDS_REMOVE_DONE
Message	SDS X (IP:Y; ID: Z) was removed successfully
Severity	Info
Description	The asynchronous process of removing an SDS has completed.
Action	None

Open SDS Device Failed

Parameter	Description
Name	OPEN_SDS_DEVICE_FAILED
Message	Failed to open a device D on SDS X (IP:Y; ID: Z) with error message E
Severity	Error
Description	Failed to open storage device D on SDS X (IP:Y; ID: Z) with error message E
Action	Check the cause of the error, and identify if it's a human error or a system malfunction. Check hardware if needed.

SDS Device Error Report

Parameter	Description
Name	SDS_DEV_ERROR_REPORT
Message	Device error reported device D on SDS X (IP:Y; ID: Z)
Severity	Error
Description	Error reported on storage device D on SDS X (IP:Y; ID: Z)
Action	Check the storage device on the server

Device capacity is high

Parameter	Description
Name	DEV_CAPACITY_USAGE_HIGH
Message	Capacity usage on %s is HIGH.")

Parameter	Description
Severity	Warning
Description	Capacity is high, due to capacity used by snapshots/thin volumes
Action	Remove unnecessary snapshots or add more storage

Device capacity is critical

Parameter	Description
Name	DEV_CAPACITY_USAGE_CRITICAL
Message	Capacity usage on %s is CRITICAL.")
Severity	Error
Description	Capacity is critical, due to capacity used by snapshots/thin volumes
Action	Remove unnecessary snapshots or add more storage

Device capacity has returned to normal

Parameter	Description
Name	DEV_CAPACITY_USAGE_NORMAL
Message	Capacity usage on %s is normal.")
Severity	Info
Description	Capacity usage is back to normal
Action	Not needed

SDS configuration has become invalid

Parameter	Description
Name	SDS_CONFIG_INVALID
Message	CLI_TARGET_NAME_CAP" %s (ID "MOS_OBJID__FORMAT ") configuration is invalid.
Severity	Critical
Description	The SDS cannot access its configuration files.
Action	Contact EMC Support

SDS disk errors were fixed

Parameter	Description
Name	SDS_FIX_DISK_ERROR
Message	CLI_TARGET_NAME_CAP" %s (ID "MOS_OBJID__FORMAT ") device %s fixed %d disk errors via reads.
Severity	Warning
Description	There were read errors on this device that were fixed by reading the data from secondary and re-writing it. This may be a sign of impending device hardware malfunction.
Action	Check for hardware malfunction

Background device scanner comparison error

Parameter	Description
Name	SCANNER_COMPARE_REPORT
Message	Background device scanner on device ID 2301536800030001 reported compare error (Device Path: <device path>, SDS: <SDS name and ID>, Peer Device Path: <peer device path>, Peer SDS: <peer SDS name and ID>, Volume name: <volume name>, Volume offset: <volume offset>)
Severity	Error
Description	Background device scanner error report, which provides details about comparison errors found during comparison of two copies of data on different devices.
Action	Check storage device for hardware malfunction

Rebuild

No Rebuild Progress Warning

Parameter	Description
Name	NO_REBUILD_PROGRESS _WARNING
Message	No rebuild progress for 30 minutes
Severity	Warning

Parameter	Description
Description	Rebuild did not progress for 30 minutes during the current recovery
Action	Contact EMC Support

No Rebuild Progress Error

Parameter	Description
Name	NO_REBUILD_PROGRESS _ERROR
Message	No rebuild progress for 60 minutes
Severity	Error
Description	Rebuild did not progress for 60 minutes during the current recovery
Action	Contact EMC Support

No Rebuild Progress Critical

Parameter	Description
Name	NO_REBUILD_PROGRESS _CRITICAL
Message	No rebuild progress for 180 minutes
Severity	Critical
Description	Rebuild did not progress for 180 minutes during the current recovery
Action	Contact EMC Support

Rebuild Progress Resumed

Parameter	Description
Name	REBUILD_PROGRESS_RESUMED
Message	Rebuild progress resumed
Severity	Info
Description	Following a detection of a rebuild not progressing, the system has now detected that the rebuild progress has resumed. The system is currently recovering.
Action	None

CHAPTER 17

Return messages

This chapter describes an overview of the messages that can be returned by the ScaleIO CLI. Topics include:

- Return messages overview 312
- Message list 312

Return messages overview

This section describes the contents and the types of return messages in the ScaleIO CLI.

Each return message contains the following:

- Name of the return message (return code)
- Message presented with the return code
- Commands that may return the relevant code
- Action or actions that may resolve the issue (if relevant)

Message list

The following section lists all ScaleIO return messages.

SDS device name changed

Parameter	Description
Name	SDS_DEVICE_NAME_CHANGED
Message	The SDS detected that the device name has changed. Re-run the command using the device ID.
Commands	`remove_sds_device`, `abort_remove_sds_device`
Description	In some situations (usually following a reboot), the operating system changes the device name. The SDS detects this change. `Remove_sds` is a time consuming operation. To ensure that you have chosen the correct device, ScaleIO alerts you to the name change.
Action	Identify the device in the output of `query_sds` and use the device ID to run the command again. You can also resynchronize the device name.

Success

Parameter	Description
Name	SUCCESS
Message	Success
Commands	All
Description	Command was executed successfully

Parameter	Description
Action	None

No resources

Parameter	Description
Name	NO_RESOURCES
Message	Insufficient resources
Commands	`add_volume`
Description	There is insufficient space to create a volume of the requested size
Action	Add capacity to the system by adding another SDS/device, or create a smaller volume

Timeout

Parameter	Description
Name	TIMEOUT
Message	A timeout occurred
Commands	All
Description	No reply received from the MDM within the specified timeout period
Action	Validate the requested operational state using queries

Already exists

Parameter	Description
Name	ALREADY_EXISTS
Message	Already Exists
Commands	All object addition commands
Description	The added object already exists
Action	Check command parameters, and use queries to view existing objects' properties

Communication error

Parameter	Description
Name	COMMUNICATION_ERROR
Message	Communication error
Commands	SDS related commands
Description	SDS has communication errors
Action	Use `query_all` to determine which SDS is disconnected

Wrong reconfiguration mode

Parameter	Description
Name	WRONG_RECONF_MODE
Message	SDS being configured, retry command
Commands	All SDS related commands
Description	SDS is being configured by the MDM
Action	Retry the command

SDS add device “open” error

Parameter	Description
Name	TGT_ADD_DEV_OPEN_ERROR
Message	Could not open SDS device
Commands	`add_sds, add_sds_device`
Description	SDS cannot open one of the specified devices
Action	Validate the command parameters. If correct, validate that the device can be accessed on the SDS.

SDS add device size problem

Parameter	Description
Name	TGT_ADD_DEV_SIZE_PROBLEM
Message	SDS device size error
Commands	`add_sds, add_sds_device`

Parameter	Description
Description	The device size is smaller or larger than the allowed size. For a full description of product limits, see the *ScaleIO User Guide*.
Action	Validate the command parameters, and set the device size appropriately

SDS add device IO error

Parameter	Description
Name	TGT_ADD_DEV_IO_ERROR
Message	Add SDS device IO error
Commands	`add_sds, add_sds_device`
Description	The SDS was unable to write to one of the specified devices
Action	Validate the command parameters, and check the specified devices

SDS actively attach to this MDM

Parameter	Description
Name	TGT_ACTIVE_ATTACH_TO_THIS_MDM
Message	SDS is actively attached to this MDM
Commands	`add_sds`
Description	SDS was already configured by this MDM. This usually implies that the provided IP address resolves to an already configured SDS CURRENTLY CONNECTED“Actively” implies that the SDS is configured and communicating with the MDM.
Action	Check the provided IP address. Contact EMC Support.

SDS already attached to this MDM

Parameter	Description
Name	TGT_ALREADY_ATTACH_TO_THIS_MDM
Message	SDS is already attached to this MDM
Commands	`add_sds`

Parameter	Description
Description	SDS was previously configured by this MDM. This usually indicates that the SDS was previously removed using `remove_sds`, but was not cleaned up properly.
Action	Validate the SDS IP address and port, and check the MDM configuration to validate that the SDS is not already configured under a different IP address. If the SDS is not found in the MDM configuration, this might indicate that it was not cleaned up properly when it was removed. To add the SDS forcefully, thus overriding its configuration and data, use the `force_clean` flag.

SDS active attach to other MDM

Parameter	Description
Name	TGT_ACTIVE_ATTACH_TO_OTHER_MDM
Message	SDS is actively attached to another MDM
Commands	`add_sds`
Description	SDS is already managed by another MDM. "Actively" implies that the SDS is configured and communicating with the MDM.
Action	Check the command parameters to validate the SDS IP address and port. Check the configuration of any other MDMs installed that might be connected to this SDS. Contact EMC Support.

SDS already attached to other MDM

Parameter	Description
Name	TGT_ALREADY_ATTACH_TO_OTHER_MDM
Message	SDS is already attached to another MDM
Commands	`add_sds`
Description	SDS was previously configured by another MDM
Action	Check your command parameters to validate the SDS IP address and port. Check for other ScaleIO installations of MDMs. If the SDS is not found in any MDM configuration, this might indicate that it was not cleaned up

Parameter	Description
	properly when it was removed. To add the SDS, use `the -force_clean` flag. **Note** This command erases the existing SDS configuration and data.

SDS not found

Parameter	Description
Name	TGT_NOT_FOUND
Message	Could not find SDS
Commands	All SDS related commands
Description	The SDS was not found
Action	Check the provided parameters. Use `query_all_sds` to view the configured SDSs.

Volume not found

Parameter	Description
Name	VOL_NOT_FOUND
Message	Could not find the volume
Commands	All volume related commands
Description	The volume was not found
Action	Check the provided parameters. Use `query_all_volumes` to view the configured volumes.

SDC not found

Parameter	Description
Name	INI_NOT_FOUND
Message	Could not find the SDC
Commands	All SDC related commands
Description	SDC was not found

Parameter	Description
Action	Use `query_all_sdc` to view the configured SDCs

Volume already mapped to SDC

Parameter	Description
Name	VOL_ALREADY_MAPPED_TO_THIS_INI
Message	The volume is already mapped to this SDC
Commands	`map_volume_to_sdc` (when mapping to a single SDC)
Description	The specified volume is already mapped to the specified SDC
Action	Check the command parameters, and use `query_all_volumes` to validate the current mapping

Volume already mapped to all SDCs

Parameter	Description
Name	VOL_ALREADY_MAPPED_TO_ALL_INIS
Message	The volume is already mapped to all SDCs
Commands	`map_volume_to_sdc` (when mapping to all SDCs)
Description	The specified volume is already mapped to all SDCs
Action	Check the command parameters, and use `query_all_volumes` to validate the current mapping

Volume not mapped to SDC

Parameter	Description
Name	VOL_NOT_MAPPED_TO_INI
Message	The volume is not mapped to SDC
Commands	`unmap_volume_from_sdc`, `query_sdc_volume_limits`, `set_sdc_volume_limits`
Description	The specified volume is not mapped to the specified SDC

Parameter	Description
Action	Check the command parameters, and use `query_all_sdc` to validate the current mapping

Volume mapped to all SDCs

Parameter	Description
Name	VOL_MAPPED_TO_ALL_INIS
Message	The volume is mapped to all SDCs
Commands	`map_volume_to_sdc`, `unmap_volume_from_sdc`, `set_sdc_volume_limits`
Description	The volume is mapped to all SDCs
Action	Check the command parameters, and use `query_all_volumes` to validate the current mapping

Too many mappings

Parameter	Description
Name	TOO_MANY_MAPPINGS
Message	There are too many mappings
Commands	`map_volume_to_sdc`
Description	Mapping of an additional volume will exceed the system limit
Action	Validate if all existing volumes that are mapped are relevant. Contact EMC Support.

Invalid volume ID

Parameter	Description
Name	INVALID_VOL_ID
Message	Invalid volume ID. Please try again with a valid ID.
Commands	All volume related commands
Description	The specified volume ID is invalid

Parameter	Description
Action	Check the specified volume ID. Use `query_all_volumes` to find the correct ID, or use the volume name.

Invalid SDS ID

Parameter	Description
Name	INVALID_TGT_ID
Message	Invalid SDS ID. Please try again with a valid ID.
Commands	All SDS commands
Description	The specified SDS ID is invalid
Action	Check the specified SDS ID. Use `query_all_sds` to find the correct ID, or use the SDS name or IP address.

Invalid SDC ID

Parameter	Description
Name	INVALID_INI_ID
Message	Invalid SDC ID. Please try again with a valid ID.
Commands	All SDC commands
Description	The specified SDC ID is invalid
Action	Check the specified SDC ID. Use `query_all_sdc` to find the correct ID, or use the SDC IP address.

Host not found

Parameter	Description
Name	HOST_NOT_FOUND
Message	Could not find the host
Commands	All SDS commands
Description	The specified host name was not found
Action	Check the specified host name

Duplicate host name

Parameter	Description
Name	DUP_HOST_NAME
Message	Duplicate host name
Commands	`add_sds`
Description	The list of host names includes two identical host names or IP addresses, or two host names are resolved to the same IP address
Action	Check the specified host name list

Address and port in use

Parameter	Description
Name	TGT_ADDR_AND_PORT_IN_USE
Message	The SDS IP address and port already in use. Please try with different ones.
Commands	`add_sds`
Description	The specified IP address and port are already in use by another defined SDS
Action	Check the specified IP address and port

Too many SDSs in Protection Domain

Parameter	Description
Name	TOO_MANY_TGTS_IN_FDOM
Message	There are too many SDSs in Protection Domain
Commands	`add_sds`
Description	Adding the SDS will exceed the system limit of allowed SDSs in a Protection Domain
Action	Check configuration

Too many SDSs

Parameter	Description
Name	TOO_MANY_TGTS
Message	There are too many SDSs

Parameter	Description
Commands	`add_sds`
Description	Adding the SDS will exceed the system limit of allowed SDSs in the system
Action	Check configuration

Too many volumes in Protection Domain

Parameter	Description
Name	TOO_MANY_VOLUMES_IN_FD
Message	Too many volumes in Protection Domain
Commands	`add_volume`
Description	Adding a volume will exceed the limit for the number of volumes allowed in a Protection Domain
Action	Check configuration

Too many volumes

Parameter	Description
Name	TOO_MANY_VOLS
Message	There are too many volumes
Commands	`add_volume`
Description	Adding a volume will exceed the limit for the number of volumes allowed in a system
Action	Check configuration

SDS name in use

Parameter	Description
Name	TGT_NAME_IN_USE
Message	SDS name in already in use. Please use a different name.
Commands	`add_sds, rename_sds`
Description	The specified SDS name is already in use by another SDS

Parameter	Description
Action	Use `query_all_sds` to validate existing names, and issue the command with another name

Volume name in use

Parameter	Description
Name	VOL_NAME_IN_USE
Message	Volume name already in use. Please use a different name.
Commands	`add_volume, rename_volume`
Description	The specified volume name is already in use by another volume
Action	Use `query_all_volumes` to validate existing names, and issue the command with another name

Duplicate SDS device name

Parameter	Description
Name	DUPLICATE_TGT_DEVICE_NAME
Message	Duplicate SDS device name
Commands	`add_sds`
Description	A device appears more than once in the `add_sds` command
Action	Validate command parameters

Duplicate SDC IP address

Parameter	Description
Name	DUPLICATE_SDC_IP
Message	More than one SDC exists for the same IP address. Use name or ID instead, or remove the duplicate SDC.
Commands	`map_volume_to_sdc, unmap_volume_from_sdc,`

Parameter	Description
	`set_sdc_volume_limits, query_sdc_volume_limits`
Description	The specified SDC IP address is not unique. It is associated with more than one SDC.
Action	Use `query_all_sdc` to find the relevant ID. Use the SDC ID in the command instead of the SDC IP address.

SDS not under removal

Parameter	Description
Name	TGT_NOT_UNDER_REMOVAL
Message	SDS is not being removed
Commands	`abort_remove_sds`
Description	The SDS is not being removed
Action	Check command parameters. If the parameters are correct, it is still possible that the SDS has already been removed.

SDS being removed

Parameter	Description
Name	TGT_BEING_REMOVED
Message	SDS is being removed
Commands	SDS related commands
Description	The SDS is currently being removed
Action	None

Volume mapped

Parameter	Description
Name	VOL_MAPPED
Message	Volume is mapped

Parameter	Description
Commands	`remove_volume, remove_consistency_group`
Description	The volume (for `remove_volume`) or any volume in the group (`remove_consistency_group`) is mapped to an SDC
Action	Unmap before removing the volume or consistency group. Use `query_all_volumes` to check current mapping.

Use MDM IP address

Parameter	Description
Name	USE_MDM_IP
Message	Use cluster MDM physical IP to access the MDM
Commands	All cluster commands
Description	The command reached the Secondary MDM. Use the physical IP address to access the MDM.
Action	Add `--mdm_ip` to the commands

Command only in cluster

Parameter	Description
Name	CMD_ONLY_IN_CLUSTER
Message	Command can only be used in cluster mode
Commands	`switch_mdm_ownership,`
Description	This command can be issued only when in cluster mode
Action	Use `query_cluster` to check the current configuration

Command only in single mode

Parameter	Description
Name	CMD_ONLY_IN_SINGLE
Message	Command can only be used in single mode
Commands	MDM cluster commands
Description	This command can be issued in Single mode only
Action	Use `switch_mdm_ownership` to switch to single mode before issuing this command

Command error when cluster degraded

Parameter	Description
Name	CMD_ERROR_CLUSTER_DEGRADED
Message	Command cannot be used when a cluster is degraded
Commands	`switch_mdm_ownership`
Description	This command cannot be issued when the cluster is degraded
Action	Check the cluster state and hardware configuration to determine the reason for cluster degradation

Too many devices

Parameter	Description
Name	TOO_MANY_DEVICES
Message	There are too many devices for SDS
Commands	`add_sds, add_sds_device`
Description	Too many devices specified (`add_sds`), or the SDS already has the maximum number of allowed devices
Action	None

SDS device not found

Parameter	Description
Name	TGT_DEVICE_NOT_FOUND
Message	Could not find the SDS device
Commands	SDS device related commands
Description	SDS device was not found
Action	Check the command parameters. Use `query_sds` to view the device details.

License too long

Parameter	Description
Name	LICENSE_TOO_LONG
Message	The license key is too long. Please check your key and enter it again.
Commands	`set_license`
Description	The specified license key is too long
Action	Validate the specified license key with the one received from EMC

License error

Parameter	Description
Name	LICENSE_ERROR
Message	License key is corrupted
Commands	`set_license`
Description	The specified license key is invalid or does not match this version. Please contact Customer Support.
Action	Validate the specified license key with the one received from EMC

Unknown license version

Parameter	Description
Name	LICENSE_UNKNOWN_VERSION
Message	Unknown license version
Commands	`set_license`
Description	The specified license key is corrupted
Action	Validate the specified license key with the one received from EMC

License not set

Parameter	Description
Name	LICENSE_NOT_SET
Message	License has not been set
Commands	`query_license, add_sds, add_sds_device, set_data_copy_limit`
Description	License was not set. Product cannot be used.
Action	Set the license before continuing to work

Configuration exceeds new license

Parameter	Description
Name	CONFIG_EXCEEDS_NEW_LICENSE
Message	The current system configuration exceeds the license entitlements
Commands	`set_license`
Description	The specified license decreases the amount of licensed capacity. The current system capacity already exceeds the specified license.
Action	Contact EMC Support to receive a license for a larger amount of capacity

License capacity exceeded

Parameter	Description
Name	CAPACITY_LICENSE_EXCEEDED
Message	The operation could not be completed. The license capacity has been exceeded.
Commands	`add_sds, add_sds_device`
Description	Adding the SDS or device exceeds the current licensed capacity
Action	Contact EMC Support to receive a license for a larger amount of capacity

License expired

Parameter	Description
Name	LICENSE_TIME_EXCEEDED
Message	The license has expired
Commands	Most commands
Description	The system time-based license has expired
Action	Contact EMC Support to receive a new license

Installation ID mismatch

Parameter	Description
Name	INSTALL_ID_MISMATCH
Message	The license installation ID does not match the ID of this system
Commands	`set_license`
Description	The provided license key does not match the current installation
Action	Contact EMC Support

Customer ID mismatch

Parameter	Description
Name	CUSTOMER_ID_MISMATCH
Message	License customer ID does not match this system
Commands	`set_license`
Description	The provided license key does not match the current installation
Action	Contact EMC Support

Unsupported license version

Parameter	Description
Name	LICENSE_UNSUPPORTED_VERSION
Message	Unsupported license version
Commands	`set_license`
Description	The provided license key does not match the current installation
Action	Contact EMC Support

Volume decrease not supported

Parameter	Description
Name	VOL_DECREASE_NOT_SUPPORTED
Message	Volume capacity can only be increased
Commands	`modify_volume_capacity`
Description	The specified volume capacity is smaller than the existing volume capacity. Decreasing volume capacity is not supported.
Action	Use `query_all_volumes` to validate your current configuration

SDS device being removed

Parameter	Description
Name	TGT_DEV_BEING_REMOVED
Message	The SDS device is being removed
Commands	`remove_sds, remove_sds_device, activate_sds_device, start_sds_device_test`
Description	The specified SDS device is currently being removed. The command cannot be executed.
Action	Wait for the device to be removed

Cannot remove last SDS device

Parameter	Description
Name	CANNOT_REMOVE_LAST_SDS_DEVICE
Message	Last remaining SDS device cannot be removed
Commands	`remove_sds_device`
Description	The command attempts to remove the last remaining SDS device. This is an illegal operation.
Action	To remove the last remaining device, remove the SDS

Device not under removal

Parameter	Description
Name	DEV_NOT_UNDER_REMOVAL
Message	The SDS device is not being removed
Commands	`abort_remove_sds`
Description	The specified device is not in the process of being removed
Action	Check the command parameters

SDS already adding device

Parameter	Description
Name	TGT_ALREADY_ADDING_DEV
Message	A device is being added to this SDS
Commands	`add_sds_device, remove_sds`
Description	This operation cannot be performed while a device is being added to the SDS
Action	Wait for the device to be added

Invalid device ID

Parameter	Description
Name	INVALID_DEV_ID
Message	The device ID is invalid. Please use a valid device ID.
Commands	SDS device related commands
Description	The specified SDS device ID is invalid
Action	Use `query_sds` to determine the correct ID

Too many Protection Domains

Parameter	Description
Name	TOO_MANY_FDs
Message	There are too many Protection Domains
Commands	`add_protection_domain`
Description	Adding a Protection Domain will exceed the system limit of the allowed Protection Domains
Action	None

Protection Domain name in use

Parameter	Description
Name	FD_NAME_IN_USE
Message	Protection Domain name already in use. Please try a different name.
Commands	`add_protection_domain, rename_protection_domain`
Description	The specified Protection Domain name is already in use by another Protection Domain
Action	Use `query_all` to validate the existing names, and issue the command with another name

Protection Domain not found

Parameter	Description
Name	FD_NOT_FOUND
Message	Could not find Protection Domain
Commands	All Protection Domain related commands
Description	The specified Protection Domain name or ID do not match any existing Protection Domains
Action	Check the command parameters. Use `query_all` to validate the existing Protection Domain names.

Protection Domain has SDSs

Parameter	Description
Name	FD_HAS_TGTS
Message	Protection Domain has SDSs
Commands	`remove_protection_domain`
Description	The Protection Domain contains SDSs and cannot be removed
Action	Remove the SDSs before attempting to remove the Protection Domain

Not enough SDSs in Storage Pool

Parameter	Description
Name	NOT_ENOUGH_TGTS_IN_STORAGE_POOL
Message	Less than three devices from different SDSs are defined in this Storage Pool. Add devices to the Storage Pool from additional SDSs.
Commands	`add_volume`
Description	The Storage Pool must contain devices from at least three different SDSs in order to create a volume
Action	Add more devices to the Storage Pool before creating a volume

Remove in progress

Parameter	Description
Name	REMOVE_IN_PROGRESS
Message	Removal of the object is in progress. Please wait.
Commands	Multiple commands
Description	The operation cannot be performed, because the volume is being removed
Action	Wait for the volume to be removed

No network test results

Parameter	Description
Name	NET_TEST_NOT_FOUND
Message	Could not find network test results
Commands	`query_sds_network_test_results`
Description	The specified SDS has no network test results
Action	Check the command parameters to make sure that a `start_sds_network_test` was sent to the specified SDS. If this is the correct

Parameter	Description
	SDS, resend the `start_sds_network_test` command.

Network test in progress

Parameter	Description
Name	NET_TEST_IN_PROGRESS
Message	Network test in progress
Commands	`query_sds_network_test_results`
Description	The SDS network test has not been completed
Action	Wait for the test to be completed. Use the same query periodically until results are displayed.

SDS device already active

Parameter	Description
Name	TGT_DEV_ALREADY_ACTIVE
Message	The SDS device is already active
Commands	`activate_sds_device`
Description	At least one specified device is already active
Action	Use `query_sds` to view the device status

SDS device wrong test mode

Parameter	Description
Name	TGT_DEV_WRONG_TEST_MODE
Message	The SDS device is in incorrect test mode
Commands	`activate_sds_device`
Description	At least one specified device will automatically be activated when its test is completed

Parameter	Description
Action	Use `query_sds` to view the device status

SDS device wrong test state

Parameter	Description
Name	TGT_DEV_WRONG_TEST_STATE
Message	The SDS device test is not complete
Commands	`activate_sds_device`
Description	At least one specified device test is not yet complete. This device cannot be activated.
Action	Use `query_sds` to view the devices' status

SDS device test in progress

Parameter	Description
Name	TGT_DEV_TEST_IN_PROGRESS
Message	The SDS device test is in progress
Commands	`start_sds_device_test`
Description	At least one specified device test is in-progress. A new test for this device cannot be started now.
Action	Use `query_sds` to view the devices' test status

Volume allocation "busy" error

Parameter	Description
Name	VOL_ALLOC_ERROR_BUSY
Message	The system is busy. Retry later.
Commands	`add_volume`
Description	The MDM is attempting to recover a connection with an SDS. You cannot create a volume now.

Parameter	Description
Action	Use `query_all` to view the system status. Retry the command when the situation has returned to normal.

Capacity highly unbalanced

Parameter	Description
Name	CAPACITY_HIGHLY_UNBALANCED
Message	System capacity is unbalanced
Commands	`add_volume`
Description	System capacity is unbalanced. This can be a temporary or permanent state. • Temporary Due to adding or removing capacity, or due to disconnections. In this case, the system will automatically work in the background to correct the situation. • Permanent The capacity of the devices added to ScaleIO is highly unbalanced. This requires user intervention to correct the situation.
Action	If a rebalance is in progress, wait for it to finish and resubmit the command. If the situation persists, contact EMC support.

Volume creation failed

Parameter	Description
Name	VOL_CREATE_FAIL
Message	Could not create the volume
Commands	`add_volume`
Description	The volume creation failed
Action	Use `query_all` to validate the system state, and retry the command

Invalid password

Parameter	Description
Name	INVALID_PASSWORD
Message	Invalid password. Please try again.
Commands	Most commands
Description	The specified password is incorrect
Action	Check that you are using the correct password in the command

No permissions

Parameter	Description
Name	NO_PERMISSIONS
Message	You do not have permission to perform this operation. Please contact your system administrator for assistance.
Commands	Most commands
Description	The specified user does not have permissions to issue this command
Action	Contact the system administrator

Too many Storage Pools in Protection Domain

Parameter	Description
Name	TOO_MANY_STORAGE_POOLS_IN_PROTECTION_DOMAIN
Message	There are too many Storage Pools in Protection Domain
Commands	`add_storage_pool`
Description	Addition of another Storage Pool will exceed the system limit of allowed Storage Pools in a Protection Domain
Action	Use `query_protection_domain` to view the Storage Pools in the Protection Domain

Storage Pool already exists

Parameter	Description
Name	STORAGE_POOL_ALREADY_EXISTS
Message	Storage Pool already exists
Commands	`add_storage_pool`
Description	A Storage Pool with the specified name already exists in the specified Protection Domain
Action	Use `query_protection_domain` to view the Storage Pools in the Protection Domain, and choose a different name

Storage Pool not found

Parameter	Description
Name	STORAGE_POOL_NOT_FOUND
Message	Could not find Storage Pool
Commands	All Storage Pool related commands
Description	The specified Storage Pool was not found
Action	Use `query_protection_domain` to view the Storage Pools in the Protection Domain, and check the issued parameters

Storage Pool has devices

Parameter	Description
Name	STORAGE_POOL_HAS_DEVS
Message	The Storage Pool has SDS devices
Commands	`remove_storage_pool`
Description	The specified Storage Pool has SDS devices associated with it
Action	Remove the SDS devices before removing the Storage Pool

Storage Pool name already exists

Parameter	Description
Name	STORAGE_POOL_NAME_ALREADY_EXISTS
Message	Storage Pool name already exists
Commands	`rename_storage_pool`
Description	A Storage Pool with the specified new name already exists in the Protection Domain
Action	Use `query_protection_domain` to view the Storage Pools in the Protection Domain. Choose another name for the Storage Pool.

Invalid Storage Pool ID

Parameter	Description
Name	INVALID_STORAGE_POOL_ID
Message	Invalid Storage Pool ID
Commands	All Storage Pool related commands
Description	The specified Storage Pool ID is invalid
Action	Check the specified Storage Pool ID. Use `query_protection_domain` to find the correct ID, or use the Storage Pool name.

Protection Domain has Storage Pools

Parameter	Description
Name	FD_HAS_STORAGE_POOLS
Message	Protection Domain has Storage Pools
Commands	`remove_protection_domain`
Description	The Protection Domain has more than one Storage Pool associated with it
Action	Remove the Storage Pools before removing the Protection Domain

SDC has mappings

Parameter	Description
Name	INI_HAS_MAPPINGS
Message	SDC has mapped volume(s)
Commands	`remove_sdc`
Description	At least one volume is mapped to the specified SDC
Action	Remove the mapping before removing the SDC

Too many volumes in Vtree

Parameter	Description
Name	TOO_MANY_VOLS_IN_VTREE
Message	There are too many volumes in the VTree
Commands	`snapshot_volume`
Description	Taking a snapshot of the specified volume will exceed the allowed system limit of snapshots per volume
Action	Use `query_volume` to view the current snapshots. Consider deleting some snapshots.

Too many snapshots in Vtree

Parameter	Description
Name	TOO_MANY_SNAPS_IN_VTREE
Message	There are too many snapshots in the VTree
Commands	`snapshot_volume`
Description	Taking a snapshot of the specified volume will exceed the allowed system limit of snapshots per volume
Action	Contact EMC Support

No volume to delete

Parameter	Description
Name	NOTHING_TO_DO
Message	No volume to delete
Commands	`remove_volume`
Description	The combination of parameters resulted in an empty set of volumes to delete
Action	Use `query_volume` to view the current state, and review the command parameters accordingly

Too many devices in Storage Pool

Parameter	Description
Name	TOO_MANY_DEVICES_IN_STORAGE_POOL
Message	There are too many devices in Storage Pool
Commands	`add_sds, add_sds_device`
Description	Adding a device to the specified Storage Pool will exceed the allowed system limit of SDS devices per Storage Pool
Action	Use `query_sds` to find the association of devices to Storage Pools

Volume not found for consistency group

Parameter	Description
Name	VOL_NOT_FOUND_FOR_SNAP_GROUP
Message	No volumes were found for the given consistency group
Commands	`remove_consistency_group`
Description	No volume found for the specified consistency group

Parameter	Description
Action	Use `query_all_volumes` to view the existing consistency groups, and review command parameters accordingly

Too many IP addresses

Parameter	Description
Name	TOO_MANY_IPS
Message	There are too many IP addresses for SDS
Commands	`add_sds_ip`
Description	Addition of an IP address to the SDS exceeds the allowed system limit of IP addresses per SDS
Action	Use `query_sds` to view the SDS IP addresses

Cannot remove last SDS IP address

Parameter	Description
Name	LAST_IP
Message	SDS must have at least one IP address
Commands	`remove_sds_ip`
Description	An SDS must have at least one IP address. Cannot remove the last remaining IP address.
Action	Use `query_sds` to view the SDS IP addresses.

SDS IP address not found

Parameter	Description
Name	SDS_IP_NOT_FOUND
Message	SDS IP not found
Commands	`remove_sds_ip`

Parameter	Description
Description	The specified SDS IP address was not found
Action	Use `query_sds` to view the SDS IP addresses

Volume size too large

Parameter	Description
Name	VOL_SIZE_TOO_LARGE
Message	The volume size too large
Commands	`add_volume, modify_volume_capacity`
Description	The specified volume size is too large
Action	Use a smaller volume size

Capacity too low for snapshot

Parameter	Description
Name	CAPACITY_TOO_LOW_FOR_SNAP
Message	The capacity of the Storage Pool is too small to create new snapshots
Commands	`snapshot_volume`
Description	Snapshots are disabled when the free capacity in the Storage Pool is very low. The threshold is defined using the `set_capacity_alerts_threshold` command. The default threshold is 10% below the predefined spare.
Action	Add storage capacity, or free up space by removing some snapshots

IP address belongs to other SDS

Parameter	Description
Name	IP_BELONGS_TO_OTHER_TGT
Message	The specified IP address belongs to another SDS. Please use the correct IP address.

Parameter	Description
Commands	`add_sds, add_sds_ip`
Description	The specified IP address is already assigned to another SDS
Action	Check command parameters

Invalid port

Parameter	Description
Name	INVALID_PORT
Message	The specified SDS port is not in the valid range. Please try again with a valid port number.
Commands	`add_sds, add_sds_ip`
Description	The specified port must be greater than 1024
Action	Check command parameters

LIA package not installed

Parameter	Description
Name	LIA_PACKAGE_NOT_INSTALLED
Message	The package is not installed.
Description	You are trying to uninstall the LIA package, but the package is not currently installed.
Action	Check why are you trying to uninstall a component that is not installed. Verify that you have the correct package name for the package that you want to uninstall.

LIA multiple packages installed

Parameter	Description
Name	LIA_MULTIPLE_PACKAGES_INSTALLED
Message	More than one package is installed.
Description	More than one package of the specified component is installed.

Parameter	Description

Action	Uninstall the packages manually:
	Linux:
	1. Find the package name:
	`rpm -qa | grep-i lia`
	The LIA package name is displayed. For example:
	`EMC-ScaleIO-lia-1.32-0.30.el6.x86_64`
	2. Uninstall the package:
	`rpm -e <PACKAGE_NAME>`
	Windows:
	• **Windows Control Panel > Add/Remove programs** > **emc-scaleio-lia**

LIA command failed

Parameter	Description
Name	LIA_COMMAND_FAILED
Message	The command failed. Please try again.
Description	This is a general error message.
Action	Look for more information about the error in the LIA log: Linux: `/opt/emc/scaleio/lia/logs` Windows: `C:\Program Files\emc\scaleio\ia\logs`

LIA invalid password

Parameter	Description
Name	LIA_INVALID_PASSWORD
Message	Token authentication was not successful
Description	The wrong LIA token was used, and authentication failed.

Parameter	Description
Action	Verify that you have the correct LIA password (token).

LIA operation not permitted

Parameter	Description
Name	LIA_OPERATION_NOT_PERMITTED
Message	The operation is not permitted by current configuration
Description	LIA was configured to disallow the operation that you tried to perform.
Action	Contact your system administrator; change the LIA configuration in your system.

LIA file is too large

Parameter	Description
Name	LIA_FILE_IS_TOO_LARGE
Message	The file is too large to send back to the user.
Description	The file returned by the `get_info` command is larger than 500 MB, and therefore cannot be sent back to you.
Action	

LIA wrong ID

Parameter	Description
Name	LIA_WRONG_ID
Message	LIA installation ID verification failed
Description	The LIA ID that was passed in the installation/uninstall package is different to the installation ID that LIA is currently using.
Action	Update LIA's installation ID by using the command `lia_msg_type__set_installation_id`

LIA disk does not exist

Parameter	Description
Name	LIA_DISK_DOES_NOT_EXIST
Message	The disk does not exist in the system
Description	The disk name that was passed with the `prepare_disk` command does not exist in the system
Action	Check that you are using the correct the disk name

SDS IP address not found

Parameter	Description
Name	TGT_IP_NOT_FOUND
Message	The SDS IP address cannot be found.
Commands	`modify_sds_ip, remove_sds_ip`
Description	You tried to remove or modify an SDS IP address that does not exist in the system.
Action	Verify that you are specifying the correct IP address in the command.

Bad threshold value

Parameter	Description
Name	BAD_THRESHOLD_VALUE
Message	The threshold value is out of range. Please enter a valid percentage value.
Commands	`set_capacity_alert_threshold`
Description	The legal range for the capacity alert threshold is a percentage value between 0 and 100. The value specified for critical alerts must be less than the value for high alerts.
Action	Specify a value between 0 and 100. Ensure that the value for critical alerts is less than the value for high alerts.

Storage Pool has Volumes

Parameter	Description
Name	STORAGE_POOL_HAS_VOLS
Message	The Storage Pool has volumes. Please remove all the volumes before removing the Storage Pool.
Commands	`remove_storage_pool` `remove_protection_domain`
Description	Only empty Storage Pools can be removed.
Action	Remove all volumes before removing the Storage Pool or Protection Domain.

Protection Domain and Storage Pool mismatch

Parameter	Description
Name	FD_AND_SP_MISMATCH
Message	The specified Storage Pool is not in the specified Protection Domain.
Commands	Any command which uses the combination of Storage Pool and Protection Domain.
Description	The operation could not be performed because the combination of Storage Pool and Protection Domain is incorrect and the Storage Pool could not be found.
Action	Use the correct combination of Storage Pool and Protection Domain in the command. Use the command `query_all` to list all the Storage Pools to find this information.

SDS IP address already exists

Parameter	Description
Name	TGT_IP_ALREADY_EXISTS
Message	This SDS IP address already exists. Please use a different IP address.
Commands	`add_sds_ip`

Parameter	Description
Description	The IP address used in the command is already in use in the system.
Action	Verify that you are using the correct IP address in the command, and that the address has not been assigned to any other SDSs.

IP role already set

Parameter	Description
Name	ROLE_ALREADY_SET
Message	The IP role for this SDS has already been configured.
Commands	`modify_sds_ip`
Description	The IP role for the specified SDS has already been configured.
Action	Verify that you are configuring the intended SDS, and that you are specifying the correct IP address.

Cluster ID mismatch

Parameter	Description
Name	CLUSTER_ID_MISMATCH
Message	The system ID number that you provided does not match the actual system ID.
Commands	`query_poll, query_properties, query_object`
Description	The system ID number has been modified, and does not match the ID specified in the command.
Action	No action required - this code is provided for information purposes for advanced users.

Device already exists

Parameter	Description
Name	DEV_ALREADY_EXISTS

Parameter	Description
Message	A device with the given name already exists in the SDS.
Commands	`add_sds_device`
Description	You tried to add a device to the system, using a name that is already defined in the SDS.
Action	Verify that the specified device has not already been added to the system.

Volume size is illegal

Parameter	Description
Name	VOL_SIZE_ILLEGAL
Message	The specified volume size is illegal. Volume size must be in 8 GB blocks.
Commands	`add_volume, modify_volume_size`
Description	You tried to define a volume size that is not a multiple of 8 GB. Volume size can only be defined in multiples of 8 GB.
Action	Specify volume size in blocks of 8 GB.

SDC must be disconnected

Parameter	Description
Name	INI_MUST_BE_DISCONNECTED
Message	The SDC must be disconnected.
Commands	`remove_sdc`
Description	Only SDCs that are currently inactive (disconnected) can be removed.
Action	This command is intended to remove old, disconnected SDCs. Do not remove active SDCs.

Invalid percentage

Parameter	Description
Name	INVALID_PERCENTAGE
Message	The percentage value that you provided is invalid. Please use a value between 0—100.
Commands	`set_thick_volume_percentage`
Description	The value that you entered was less than 0 or greater than 100.
Action	Use a value between 1 and 100 in the command.

Wrong receive group

Parameter	Description
Name	WRONG_RECEIVE_GROUP
Message	The command was sent to the wrong receive group due to component software mismatch.
Commands	All commands
Description	There are different software versions installed on some components in your system, and this is causing a software mismatch error.
Action	Verify that the same software version is installed on all components, and fix, if necessary.

Mismatched SW ID

Parameter	Description
Name	MISMATCHING_SWID
Message	The license contains a mismatch of the SWID number. Please contact Customer Support.
Commands	`set_license`
Description	The SWID number that you entered does not match your license.
Action	Contact Customer Support

Invalid license issuer

Parameter	Description
Name	INVALID_LICENSE_ISSUER
Message	The issuer of the license you are attempting to add does not match that of the product.
Commands	`set_license`
Description	There is a problem with your license.
Action	Contact Customer Support

Number of parallel messages is too low

Parameter	Description
Name	NUM_PARALLEL_MSG_TOO_LOW
Message	The number of parallel messages is too low. Please use a number in the range 1 - 6.
Commands	`start_sds_network_test`
Description	You used a number less than 1 in the command.
Action	Use a number in the range of 1 to 6 for parallel messages in the command.

Number of parallel messages is too high

Parameter	Description
Name	NUM_PARALLEL_MSG_TOO_HIGH
Message	The number of parallel messages is too high. Please use a number in the range 1 - 6.
Commands	`start_sds_network_test`
Description	You used a number greater than 6 in the command.
Action	Use a number in the range of 1 to 6for parallel messages in the command.

Network test size is too high

Parameter	Description
Name	NETWORK_TEST_SIZE_TOO_HIGH
Message	The network test size is too high. Please use a test size less than 10.
Commands	`start_sds_network_test`
Description	You used a number greater than 10 in the command.
Action	Use a number less than 10 for test size in the command.

Network test length too long

Parameter	Description
Name	NETWORK_TEST_LENGTH_TOO_HIGH
Message	The network test length is too long. Please use a test length less than 43200 seconds.
Commands	`start_sds_network_test`
Description	You used a number greater than 43200 for test length in the command.
Action	Use a number less than 43200 for test length in the command.

Network test must limit either time or size

Parameter	Description
Name	NETWORK_TEST_MUST_LIMIT_EITHER_TIME_OR_SIZE
Message	The network test must limit either time or size.
Commands	`start_sds_network_test`
Description	Your command did not contain a value for time or size.

Parameter	Description
Action	Include a value for either time or size in the command.

Not enough SDSs in Protection Domain for test

Parameter	Description
Name	NOT_ENOUGH_TGTS_IN_FD_FOR_TEST
Message	To start the test, there must be at least 2 SDSs in the Protection Domain.
Commands	`start_sds_network_test`
Description	There are less than two SDSs in the Protection Domain, and therefore the test operation cannot be performed.
Action	Ensure that there are at least 2 SDSs in the Protection Domain before attempting to test it.

Not allowed in Storage Pool with devices

Parameter	Description
Name	NOT_ALLOWED_IN_SP_WITH_DEVS
Message	This operation is only allowed when there are no devices in the Storage Pool. Please remove all devices from the Storage Pool.
Commands	`modify_zero_padding_policy`
Description	You can only modify the zero padding policy of empty Storage Pools.
Action	Ensure that all devices are removed from the Storage Pool before you attempt to modify its zero padding policy.

License capacity mismatch

Parameter	Description
Name	LICENSE_CAPACITY_MISMATCH

Parameter	Description
Message	The license contains a mismatch of the capacity values for basic and advanced features. Please contact Customer Support.
Commands	`set_license`
Description	The capacity allowed for basic and advanced features in your license do not match your installation.
Action	Contact Customer Support.

Exceeds system capacity limitations

Parameter	Description
Name	EXCEED_SYSTEM_CAPACITY_LIMITATIONS
Message	The request exceeds system limitations
Commands	`add_volume, modify_volume_capacity`
Description	The command specifies a volume larger in size than the maximum system capacity allowed by the system.
Action	Remove one or more volumes in the system, to make room for the new one.

SDC mapped to too many volumes

Parameter	Description
Name	INI_MAPPED_TO_TOO_MANY_VOLS
Message	The SDC is mapped to too many volumes. Please remove some mapping.
Commands	`map_volume_to_sdc`
Description	You attempted to map more than the maximum allowed number of volumes (8192) to one or more SDCs.
Action	Avoid mapping more than the maximum allowed volumes to a single SDC.

Not enough devices in Storage Pool for removal

Parameter	Description
Name	NOT_ENOUGH_DEVICES_IN_STORAGE_POOL_FOR_REMOVAL
Message	Not enough devices in Storage Pool.
Commands	`remove_sds, remove_sds_device`
Description	Each Storage Pool must contain at least two devices. If there are two or less devices in the Storage Pool and you try to remove one, the command will fail.
Action	Add devices to the Storage Pool before trying to remove this one.

Invalid network limits

Parameter	Description
Name	INVALID_TGT_NETWORK_LIMITS
Message	Bad limits given. Overall must be larger than both rebuild and rebalance and all limits must be over 5MB
Commands	`sds_network_limits`
Description	The total (overall) limit must be greater than the sum of the Rebuild limit and the Rebalance limit. All limits must be greater than 5 MB.
Action	Ensure that limit values are greater than 5 MB, and that the overall limit value is greater than Rebuild limit+Rebalance limit.

A device is currently being removed

Parameter	Description
Name	ONE_TGT_DEV_BEING_REMOVED
Message	This command cannot be completed; An SDS device is currently being removed.

Parameter	Description
Commands	`activate_sds_device` `activate_all_devices`
Description	If a device is in the process of being removed, you cannot use the `activate_all_devices` command to activate the remaining ones.
Action	Wait until the device removal is complete, and try again.

A device is already active

Parameter	Description
Name	ONE_TGT_DEV_ALREADY_ACTIVE
Message	This command cannot be completed; An SDS device is currently active.
Commands	`activate_sds_device` `activate_all_devices`
Description	A device is already active.
Action	Do not use the `activate_all_devices` option in the command. Instead, specify each device individually in the command.

RAM Cache wrong state

Parameter	Description
Name	RMCACHE_WRONG_STATE
Message	Wrong RAM Cache state

Commands	`enable_sds_rmcache`
	`disable_sds_rmcache`
	`set_sds_rmcache_size`

Description	The cache is in transient state, so the requested operation cannot be executed.
Action	Wait for a few seconds and try again.

More than one mapping not allowed

Parameter	Description
Name	NOT_ALLOWED_MORE_THAN_ONE_MAPPING
Message	Only a single SDC may be mapped to this volume at a time
Commands	`map_volume_to_sdc`
Description	The volume is already mapped, and the `allow_multi_map` flag was not used in the command
Action	The volume is already mapped. Use the `allow_multi_map` flag to map the volume to additional SDCs.

At least one SDS failed

Parameter	Description
Name	AT_LEAST_ONE_TGT_FAILED
Message	At least one SDS failed

Commands	`enable_sds_rmcache`
	`disable_sds_rmcache`
	`set_sds_rmcache_size`
	for the entire Protection Domain

Description	You tried to enable/disable or change size of cache for all SDSs in the Protection Domain, and this failed for at least one of the SDSs
Action	Check the state of cache in all the SDSs, and try again for the specific SDS with the problem.

A device is in incorrect test state

Parameter	Description
Name	ONE_TGT_DEV_WRONG_TEST_STATE
Message	This command cannot be completed; An SDS device is currently in incorrect test mode

Parameter	Description
Commands	`activate_sds_device` `activate_all_devices`
Description	A device was not added with the `test_only` option, so it cannot be activated with this command.
Action	It will activate itself on its own. No action required.

A device test is in progress

Parameter	Description
Name	ONE_TGT_DEV_WRONG_TEST_MODE
Message	This command cannot be completed; An SDS device test is currently in progress
Commands	`activate_sds_device` `activate_all_devices`
Description	A device test is currently in progress, and the device cannot be activated until the test is finished.
Action	Wait for the device test to finish, and then try again.

Enterprise features not enabled

Parameter	Description
Name	ENTERPRISE_FEATURES_NOT_ENABLED
Message	Enterprise features are not enabled for this system.
Commands	Any command that requires a special license
Description	The command that you tried to use cannot be used with the type of license that you currently have activated in your system.
Action	Verify that you are using a license with enterprise features enabled. If problems persist, contact Customer Support.

Volume already mapped to an SDC

Parameter	Description
Name	VOL_ALREADY_MAPPED_TO_AN_INI
Message	The volume is already mapped to a SDC
Commands	`map_volume_to_sdc` with the `all_sdcs` option
Description	If the volume is already mapped to one or more SDCs, you cannot use the `all_sdcs` option.
Action	Remove the volume's mapping to the SDCs, and then try again.

SDS already exists and is being removed

Parameter	Description
Name	TGT_ALREADY_EXISTS_AND_BEING_REMOVED
Message	The SDS already exists and is being removed. Please wait for the operation to finish.
Commands	`add_sds`
Description	If you try to add an SDS that already exists, and is in the process of being removed, this operation will fail. **Note** If an `add_sds` command failed, the system will try to remove it automatically, without user intervention. Therefore, if you issue the same `add_sds` command twice in a row, and the first time failed, this return code may be generated (depending on why adding the SDS failed).
Action	Verify that you are adding the correct SDS. If you are adding an SDS which is in the process of being removed, wait for the removal process to finish before trying to add the SDS back to the system.

Invalid spare percentage

Parameter	Description
Name	INVALID_SPARE_PERCENTAGE
Message	The spare percentage value that you provided is invalid. Please use a value between 0 - 99.
Commands	`modify_spare_policy`
Description	The spare percentage required in the command is invalid
Action	Verify that correct input parameters (a value between 0-99) are used with the command

Invalid background device scanner bandwidth limit

Parameter	Description
Name	INVALID_SCANNER_BW_LIMIT
Message	The scanner bandwidth limit is invalid. Please use a value between 10KB - 1GB
Commands	`enable_background_device_scanner`
Description	The bandwidth value that you provided for the bandwidth limit parameter is invalid.
Action	Verify that correct input parameters (a value between 10KB–1GB) are used with the command

Zero padding with background device scanner data comparison mode

Parameter	Description
Name	ZEROPAD_AND_SCANNER
Message	Zero padding must be enabled in order to set scanner to Data Comparison mode
Commands	`enable_background_device_scanner`
Description	The background device scanner has two modes: device only, and data comparison mode. The Storage Pool must have zero

Parameter	Description
	padding enabled before you attempt to enable background device scanning in data comparison mode.
Action	Enable zero padding in the Storage Pool (you must do this before adding devices to the Storage Pool)

IP address already assigned to another SDS

Parameter	Description
Name	ADDRESS_BELONGS_TO_DIFFERENT_TGT
Message	The given address belongs to a different SDS. Please check the system configuration.
Commands	`add_sds, add_sds_ip`
Description	This message is returned if the specified IP address has already been assigned to another SDS in the system
Action	Check SDS network configuration and IP address assignments

Capacity too low for thin volumes

Parameter	Description
Name	CAPACITY_TOO_LOW_FOR_THIN_VOL
Message	Storage Pool reached critical capacity utilization. Unable to create new thin volumes.
Commands	`add_volume, modify_volume_capacity`
Description	Thin volumes cannot be allocated if the specified Storage Pool has reached the critical capacity level
Action	Add more devices to the Storage Pool

Too many Storage Pools in the system

Parameter	Description
Name	TOO_MANY_STORAGE_POOLS

Parameter	Description
Message	There are too many Storage Pools in the system
Commands	`add_storage_pool`
Description	Addition of another Storage Pool will exceed the system limit of allowed Storage Pools
Action	Remove unused Storage Pools. If the problem is not solved, install a new system.

Remote access to the MDM is blocked

Parameter	Description
Name	REMOTE_PERMISSION_DENIED
Message	Permission denied. Remote read-only limit state is enabled and is blocking this operation from a remote location. Run this command from the Master MDM host.
Commands	All
Description	Remote read-only limit state is enabled and is blocking this operation from a remote location
Action	Run the command from the Master MDM host

Cannot add a partitioned device

Parameter	Description
Name	TGT_ADD_DEV_IS_PARTITIONED
Message	Unable to add a device that was already partitioned
Commands	`add_sds, add_sds_device`
Description	The added device is partitioned and cannot be used by the system
Action	Remove partitions from the device, or add the partitioned device to the system

Cannot add a mounted device

Parameter	Description
Name	TGT_ADD_DEV_IS_MOUNTED
Message	Unable to add a device that was already mounted
Commands	`add_sds, add_sds_device`
Description	The added device is mounted and cannot be used by the system
Action	Unmount the device

Device header mismatch

Parameter	Description
Name	INVALID_DEVICE_HEADER_SIGNATURE
Message	Found an invalid device header signature
Commands	`add_sds, add_sds_device`
Description	The device was previously added to the system, or after SDS reboot the signature of the header does not match the expected signature
Action	Check the specific SDS devices configuration

CHAPTER 18

ScaleIO on Xen

This chapter describes topics that relate to using ScaleIO on XEN. Topics include:

- Overview of ScaleIO on Xen 368
- Adding a volume in XenServer environment 368
- Removing a ScaleIO volume from Xen 369
- Modifying the size of a ScaleIO volume 370
- Xen v6.5 High Availability 371

Overview of ScaleIO on Xen

This section describes an overview of using ScaleIO on XEN.

ScaleIO best practice is to install all ScaleIO components on dom0. The installation and configuration of ScaleIO objects are the same as in a regular Linux system.

By default, dom0 comes with approximately 800 MB of memory. This might not be enough if an SDS and MDM are installed together.

It is recommended to increase dom0 memory to 4 GB. For details of how to do this, see http://support.citrix.com/article/CTX134951.

This appendix contains additional commands that must be performed on the hypervisor when adding or removing a ScaleIO volume or changing its volume.

Note

In Xen, RAM read cache is limited to 1GB.

In the Xen environment, all ScaleIO CLI commands begin with `siocli`, and not `scli`.

Adding a volume in XenServer environment

You enable use of volumes in HA, by enabling the volumes to be recognized as HBA.

This procedure is relevant for XenServer v6.5 and XenServer v7.0.

Procedure

1. Use the ScaleIO CLI to add and map a ScaleIO volume, as described in "Creating volumes" in the *ScaleIO User Guide*.
2. Get the host UUID by running the following command:

```
xe host-list
```

3. For XenServer v6.5 and v7.0, edit the file `/etc/lvm/lvm.conf` by editing the lines that starts with `types`, and adding `"scini", 16` inside the square brackets.

   Example:

```
types = ["nvme", 64, "mtip32xx", 64, "scini", 16]
```

4. For XenServer v7.0 (only), edit `/etc/lvm/master/lvm.conf`, as follows:
   a. Locate the lines that begin with `types`.
   b. In each of these lines, add this string inside the square brackets: `"scini", 16`

```
types = ["nvme", 64, "mtip32xx", 64, "scini", 16]
```

5. Use the retrieved host UUID while running the `sr-create` command.

Note

ScaleIO provides a unique ID to each volume. It is highly recommended to use the unique ID when running on XenServer. For example, the ScaleIO volume name in the hypervisor is `/dev/disk/by-id/scsi-emc-vol-4a7987a751237ae0-3d467d3900000000`.

Example

```
xe sr-create host-uuid=09fa5d27-aa08-4c71-86bb-71dc73e9f59f
content-type="ScaleIO" name-label="ScaleIO" shared=true
device-config:SCSIid=emc-
vol-4a7987a751237ae0-3d467d3900000000 type=lvmohba
```

Note

To add a shared storage repository, the following conditions must be fulfilled:

- All nodes in the XenServer Center Storage Pool must be installed with SDC.
- The ScaleIO volume to be used as the shared SR must be mapped to all SDCs in the Storage Pool.

Removing a ScaleIO volume from Xen

There are steps to perform before removing a ScaleIO volume in a Xen environment.

Before unmapping a volume, perform the following:

Procedure

1. Get the SR UUID:

```
xe sr-list
```

2. Get the PB UUID:

```
xe pbd-list sr-uuid=<sr_uuid>
```

where *<sr_uuid>* is the result from the previous step

3. Unplug the PBD:

```
xe pbd-unplug uuid=<pbd_uuid>
```

where *<pbd_uuid>* is the result from the previous step

4. Destroy the PBD:

```
xe pbd-destroy uuid=<pbd_uuid>
```

5. Forget the PBD:

```
sr-forget uuid=<sr_uuid>
```

where *<sr_uuid>* is the result from the first step

Example:

```
xe sr-list
xe pbd-list sr-uuid=4232efb0-7610-b18f-51ee-46bf377021d2
xe pbd-unplug uuid=c478e01f-eb5a-237f-9ed3-9c1c9173431b
xe pbd-destroy uuid=c478e01f-eb5a-237f-9ed3-9c1c9173431b
xe sr-forget uuid=4232efb0-7610-b18f-51ee-46bf377021d2>
```

Results

You can now use the SCLI to remove a volume. To run CLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *ScaleIO CLI Reference Guide*.

Modifying the size of a ScaleIO volume

Procedure

1. Use the CLI to modify the ScaleIO volume size.
2. Use the Xen command line to modify the volume size.

   In the following example the volume name in Dom0 is `/dev/disk/by-id/scsi-emc-vol-593b29a1640c4d79-0563e40f00000000`

   Example:
3. Issue the following command in the Xen command line.
   - Xen 6.5:

     ```
     pvresize /dev/disk/by-id/scsi-emc-
     vol-593b29a1640c4d79-0563e40f00000000
     ```
   - Xen 7.0:

     ```
     pvresize /dev/disk/by-id/scsi-emc-
     vol-593b29a1640c4d79-0563e40f00000000 --config
     global{metadata_read_only=0}
     ```
4. Check the output to validate that the new size is set:

   ```
   pvs
   ```
5. Read the UUID from `xe sr-list`.

6. In the Xen command line, issue the following command with the UUID that was obtained with `xe sr-list` as *<UUID>*.

```
xe sr-scan uuid=<UUID>
```

Xen v6.5 High Availability

Enable use of volumes in HA, by enabling the volumes to be recognized as HBA.

Procedure

1. Install ScaleIO as described in the *EMC ScaleIO Deployment Guide*.
2. Modify the `lvm.conf` file, as described in the "Adding a volume" section of the *ScaleIO User Guide*, Xen appendix.
3. Create and map a ScaleIO volume to your Xen hosts.
4. From the Xen pool master, list the available storage:

```
#ls -l /dev/disk/by-id/
```

This will return an overview of the attached ScaleIO volumes, similar to the following:

```
# scsi-"volume id" -> ../../scinia
```

5. From the Xen pool master, issue the Xen Storage Repository create command:

```
# xe sr-create name-label="Any_name" content-type="ScaleIO"
shared=true device-config:SCSIid="volume id" type=lvmohba
```

where `volume id` is the value of volume id identified in the `scsi-"volume id"` output in the previous step.

Results

Your ScaleIO volume should now appear on your Xen hosts as a Storage Repository (SR).

CHAPTER 19

Configuring ScaleIO in OpenStack Environments

This chapter contains information about ScaleIO provisioning in an OpenStack cloud operating system environment. Topics include:

- Overview 374
- Volume limitations 374
- Supported operations 374
- QoS functionality 374
- Thin provisioning 375
- Multiple Protection Domains and Storage Pools 375
- Pool-aware scheduler support 377

Overview

OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a data center. These are all managed using a dashboard or a command line interface that gives administrators control, while allowing their users to provision resources through a web-based interface. OpenStack is used in a wide variety of industries and use cases, and is supported by more than 400 leading IT hardware and software companies, who have contributed to development of the OpenStack project.

In the case of Block storage, OpenStack provides the Cinder solution, which is a block storage solution for use with servers and applications. Cinder is designed to work with widely available virtualization technologies, bare metal, and high-performance computing configurations. It can integrate with legacy systems and with third-party technologies, such as EMC ScaleIO.

Volume limitations

ScaleIO volume size is limited to a basic granularity of 8 GB. If volume size is not a multiple of 8, the size is rounded up. For example: a request to create a volume of 10 GB will create a volume of 16 GB. In OpenStack, the volume will be displayed with its requested size (10) instead of its actual size (16). A corresponding warning message will be printed to the Cinder log.

The Nova compute database retains a “connection info” field which contains the details of each OpenStack volume. If the ScaleIO Gateway’s IP address and port number are modified, these details are not updated in the Nova compute database. Therefore, do not change the ScaleIO Gateway's IP address and port number when there are active instances of ScaleIO volumes.

Supported operations

The ScaleIO Cinder driver supports the following operations:

- Create volume
- Delete volume
- Attach volume
- Detach volume
- Create snapshot
- Delete snapshot
- Create volume from snapshot
- Copy image to volume
- Copy volume to image
- Extend volume

QoS functionality

The ScaleIO Cinder driver allows limiting the IOPS and bandwidth that one SDC generates for a volume. This is done with Cinder Volume Type and Extra Specs. The

following keys should be added to the volume type to set the SDC's limits for volumes of this type:

- `sio:iops_limit` – limits the volume IOPS. The number of IOPS must be larger than 10. The `0` value represents unlimited IOPS.
- `sio:bandwidth_limit` – limits the volume network bandwidth. The bandwidth is in KB/Sec. The `0` value represents unlimited bandwidth.

Because the limits are per SDC, they will be applied after the volume is attached to an instance, and thus to a compute node/SDC.

Thin provisioning

The Cinder driver supports creation of thin provisioned volumes, in addition to thick provisioning. The provisioning type settings should be added as an extra specification of the volume type, as follows:

`sio:provisioning_type` – `thick/thin`

If the provisioning type is not specified, the default value of `thick` will be used.

Multiple Protection Domains and Storage Pools

The ScaleIO OpenStack plug-in supports two modes for working with user defined Storage Pools and Protection Domains: static configuration mode, and dynamic mode.

Static configuration mode

Static configuration mode only supports a single Storage Pool and Protection Domain. Use the `protection_domain_name` or `protection_domain_id` and `storage_pool_name` or `storage_pool_id` in the driver configuration file `cinder_scaleio.config`. When the Storage Pool name or ID is omitted, the Default Storage Pool is used.

Dynamic mode

Dynamic mode lets you work with multiple Storage Pools and Protection Domains.

Configuring dynamic mode

Procedure

1. Create volume types using either the Horizon dashboard or the CLI.
2. Using the CLI set command, assign extra specifications to each volume type, to give it some meaningful data, which is basically meta data that is assigned to the volume.

Note

Volume Type is an abstraction of various properties of a volume. Everything meaningful about the volume type is stored as a key-value pair in the `volume_type_extra_spec` table and can be viewed using the `cinder extra-specs-list` command. The ScaleIO Cinder driver lets you choose a specific Protection Domain and a Storage Pool by using the meta data of the volume during creation.

Dynamic mode example using the CLI

The following example shows the creation of two volume types: High Performance and Standard Performance.

For High Performance, `sio:sp_name` is set to `Performance_Storage` and `sio:pd_name` is set to `common`.

Similarly, for Standard Performance, `sio:sp_name` is set to `Capacity_Storage` and `sio:pd_name` is set to `common`.

Then, an 8GB volume is created with the volume type `High_Performance`, which means it will be created in the `common` Protection Domain, over the Storage Pool `Performance_Storage`.

```
cinder --os-username admin --os-tenant-name admin type-create
"High_Performance"
cinder --os-username admin --os-tenant-name admin type-create
"Standard_Performance"
cinder --os-username admin --os-tenant-name admin type-key
"High_Performance" set sio:sp_name=Performance_Storage
cinder --os-username admin --os-tenant-name admin type-key
"High_Performance" set sio:pd_name=common
cinder --os-username admin --os-tenant-name admin type-key
"Standard_Performance" set sio:sp_name=Capacity_Storage
cinder --os-username admin --os-tenant-name admin type-key
"Standard_Performance" set sio:pd_name=common
cinder --os-username admin --os-tenant-name admin create -
volume_type High_Performance -display_name test_multi 8
```

Example notes

1. Either `sio:sp_name` or `sio:sp_id` should be set, but not both. The same principle applies for Protection Domain: either `sio:pd_name` or `sio:pd_id` should be set, but not both.
2. Volume Type names `High_Performance` and `Standard Performance` are user-defined and can hold any arbitrary strings. Extra spec keys `sio:pd_id`, `sio:pd_name`, `sio:sp_id`, and `sio:sp_name` are reserved strings, and must appear exactly as shown in this example.
3. The extra spec value `Capacity_Storage` is your ScaleIO Storage Pool name. The extra spec value `common` is your ScaleIO Protection Domain name.

Defaults and missing arguments

Cinder's `VolumeType` is used to define a Storage Pool and Protection Domain, and to override the Protection Domain and Storage Pool settings specified in the configuration file. The driver will search for `volume_type` first. If `volume_type` is specified when creating a volume, the driver will search for the `volume_type` definition, and will find the matching Storage Pool and Protection Domain. If `volume_type` is not specified, the driver will fall back to use the Storage Pool and Protection Domain defined in `cinder_scaleio.config`. The settings defined in the configuration file will also be used as a fall back if the requested `volume_type` has no Storage Pool or Protection Domain definition.

Pool-aware scheduler support

From the Juno release and onwards, the Cinder driver was enhanced to support the new OpenStack feature, the pool-aware scheduler. The driver now reports capacities per pool, and not for the whole system, as in previous versions.

The pools that are reported are user-defined pools in the new configuration parameter `storage_pools`, in the `/etc/cinder/cinder.config` file.

To configure the pool-aware scheduler, perform the following steps:

Procedure

1. Use a text editor to edit lines in the `/etc/cinder/cinder.config` file.
2. For the `storage_pools` parameter, add comma-separated pairs of domain and pool that you want to work with in OpenStack.

   The format of the pair must be: *<domain_name>*:*<pool_name>*.

   For example:

   ```
   storage_pools=Lab:default, Lab2:flash
   ```

CHAPTER 20

REST API Reference

This chapter describes the REST API exposed by ScaleIO. You can use this API to manage your ScaleIO system.

- Before you begin 380
- CLI/REST API command reference 384
- Software type resources 390
- Responses 516
- Changes that might require adjustments of old code 517
- Using the IM REST API 519

Before you begin

The REST API is served from the ScaleIO Gateway (which includes the REST gateway). The ScaleIO Gateway can be installed on any operating system that is supported by ScaleIO for its core components (MDM/SDS/SDC/etc.). The ScaleIO Gateway connects to a single MDM and serves requests by querying the MDM and reformatting the answers it receives from the MDM in a RESTful manner, back to a REST client. The gateway is stateless. It requires the MDM user name and password for the login request. The login returns a token in the response, that is used for later authentications for other requests. Every ScaleIO SCLI command is also available in the ScaleIO REST API. The responses returned by the ScaleIO Gateway are formatted in JSON format.

The API is available as part of the ScaleIO Gateway package. If you used the ScaleIO Installation Manager to install ScaleIO, the gateway has already been installed and configured with the MDM details.

To install the gateway manually, see the *ScaleIO Deployment Guide*.

Before you start using this API:

1. If you installed the ScaleIO Gateway manually, ensure that the MDM connection properties are configured correctly in the ScaleIO Gateway configuration file. For more information, see "Configuring connection properties in the ScaleIO Gateway". If the Installation Manager was used to install ScaleIO, the ScaleIO Gateway is already configured, and no action is required.

   Note

   You can also run this test from the link in the Installation Manager Maintenance Operation screen.

2. Run a self-test to verify connectivity and correct configuration via a browser by navigating to the test page at this URL: https://<REST_SERVER>/rest.jsp. In addition, there is a link from the IM-WEB maintenance page (`maintain.jsp`) to the REST test page (`rest.jsp`).

   The remainder of this section covers the following topics:

Configuring connection properties in the ScaleIO Gateway

This section describes how to configure MDM connection properties and passwords in the ScaleIO Gateway. If the Installation Manager was used to install ScaleIO, the ScaleIO Gateway is already configured, and there is no need to perform these steps.

Note

A ScaleIO Gateway deployed on RHEL6 requires a specific version of NSS. For more information, see "Approved encryption methods" in the *ScaleIO User Guide*.

Configuring the Gateway using REST API commands

If the ScaleIO Gateway was installed manually after ScaleIO was installed, you can use the following REST API commands to configure it.

Note

In order to configure the ScaleIO Gateway, you must authenticate first. For more information, see General information on page 390.

When connecting to the ScaleIO Gateway to configure it, use the following connection properties: *User*: admin *Password*: *<GATEWAY_ADMIN_PASSWORD>*. When connecting to the ScaleIO Gateway to configure ScaleIO, use the following connection properties: *User*: admin *Password*:*<MDM_PASSWORD>*.

Procedure

1. Read (GET) the configuration:

```
HTTPS://<server_ip>:443/api/Configuration
```

2. Set (POST) the configuration:

```
HTTPS://<server_ip>:443/api/updateConfiguration
```

```
{
"mdmAddresses":["<MDM_MANAGEMENT_IP_ADDRESS1>",
"<MDM_MANAGEMENT_IP_ADDRESS2>"],
"mdmPort":"<MDM_PORT>",
"gatewayAdminPassword":"<ADMIN_PASSWORD>",
}
```

For example:

```
HTTPS://10.76.61.100:443/api/updateConfiguration
```

```
{
"mdmAddresses":["10.76.61.32", "10.76.61.33"],
"mdmPort":"6611",
"gatewayAdminPassword":"123",
}
```

Configuration is complete.

Using REST API to create the MDM cluster

Post `/api/instances/System/action/createMdmCluster` with a body that contains the property `updateConfiguration true` and leave the `mdmAddresses` configuration empty.

The `mdmAddresses` field will be populated with the MDM IP addresses according to the request. `createMdmCluster` creates a new user called "admin" with the password `admin`. This password must be changed before you continue working with the MDM cluster, as follows:

```
POST /api/instances/User/action/setPassword
```

```
Body:
```

```
{"oldPassword":"admin", "newPassword":<new password>}
```

Note

The `addStandbyMdm` parameter does not add the IP addresses to the `mdmAddresses` field. They must be updated manually. For more information, see "Examples for updating the mdmAddresses field".

Examples for updating the mdmAddresses field

POST `/api/updateConfiguration` body contains `mdmAddresses` with comma-separated MDM IP addresses. For example:

```
{"mdmAddresses":["10.76.60.150", "10.76.60.11">
```

You can also update this field manually, as described in "Configuring the Gateway by editing the user properties file".

Configuring the Gateway by editing the user properties file

If the ScaleIO Gateway was installed manually after ScaleIO was installed, you can configure the following Installation Manager properties, by editing the `gatewayUser.properties` file:

- Enable Gateway (default: true)

  To disable, set `features.enable_gateway=false`. You can completely disable the use of the default port, 443, by setting both this property and `features.enable_IM` to `false`.

- Change MDM port (default=6611)

  To change the port, set `mdm.port=`*<MDM_PORT>*

- Change MDM IP addresses

  To change the MDM IP addresses, set `mdm.ip.addresses=`*<MDM_MANAGEMENT_IP_ADDRESS1>; <MDM_MANAGEMENT_IP_ADDRESS1>*

  Note

  The `gateway-admin.password` property cannot be set manually, because it is saved in the file in hashed format. You should configure it using a REST API command, as explained in "General information".

  To edit the properties, perform the following:

Procedure

1. Use a text editor to open the `gatewayUser.properties` file, located in the following directory on the Installation Manager/Gateway server:

Gateway installed on	Location of gatewayUser.properties file
Windows	C:\Program Files\EMC\ScaleIO\Gateway\webapps\ROOT\WEB-INF\classes\
Linux	/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes

2. Edit the file with the desired changes.
3. Save and close the file.
4. Restart the scaleio-gateway service:
   - Windows: Restart the EMC ScaleIO Gateway service.
   - Linux: Enter the following command:

   ```
   service scaleio-gateway restart
   ```

   Configuration is complete.

Configuring the Gateway using a CLI tool

This section provides an example of gateway configuration using the CURL CLI tool.

Note

This action requires a token (key), that is received in an authentication response. For more information, see General information on page 390.

Procedure

1. Send a login request:

   ```
   curl -k -v --basic --user admin:<scaleio admin password>
   https://<host address>/api/login
   ```

   where *<scaleio admin password>* is the password used to login to the ScaleIO CLI (scli). The output of this command should be used in the following step instead of `<key>`.

2. Send the following command:

   ```
   curl -k -v --basic --user admin:<key received at login> -X
   POST -H "Content-Type: application/json" -d '{"mdmAddresses":
   ["<address a>", "<address b>"], "mdmPort":"6611",
   "gatewayAdminPassword":"<gateway admin new password>"}'
   https://<host address>/api/updateConfiguration
   ```

   Example:

   ```
   curl -k -v --basic --user
   admin:YWRtaW46MTQyNzczMTA3MzExNjplNzM2MDY5NjhlNjMyZDdjZDU4MTk3
   MTZkNWYwNzI4Mg -X POST -H "Content-Type: application/json" -d
   '{"mdmAddresses":["10.76.60.1", "10.76.60.2"],
   "mdmPort":"6611", "gatewayAdminPassword":"NewPassword"}'
   https://gateway.scaleio.com/api/updateConfiguration
   ```

   where:

   - the `-k` flag is for skipping the SSL certificate verification
   - the `-v` is an optional flag for extra output verbosity

Logs

Logs for REST activity are saved in the following locations on the ScaleIO Gateway:

Windows:

`C:\Program Files\EMC\ScaleIO\Gateway\logs\api_operations.log`

Linux:

`/opt/emc/scaleio/gateway/logs/api_operations.log`

CLI/REST API command reference

Refer to this table to determine which REST API URIs are used for specific CLI commands.

Table 32 ScaleIO CLI commands

ScaleIO CLI command	REST API URI
--abort_remove_sds	/api/instances/Sds::{id}/action/abortRemoveSds
--abort_remove_sds_device	/api/instances/Device::{id}/action/abortRemoveDevice
--abort_upgrade	/api/instances/System/action/abortUpgrade
--activate_protection_domain	/api/instances/ProtectionDomain::{id}/action/activateProtectionDomain
--activate_sds_device	/api/instances/Device::{id}/action/activateDevice
--add_certificate	/api/trustHostCertificate/Mdm
--add_fault_set	/api/types/FaultSet/instances
--add_ldap_service	/api/instances/System/action/addLdapService
--add_protection_domain	/api/types/ProtectionDomain/instances
--add_sdc	/api/types/Sdc/instances
--add_sds	/api/types/Sds/instances
--add_sds_device	/api/types/Device/instances
--add_sds_ip	/api/instances/Sds::{id}/action/removeSdsIp
--add_sds_rfcache_device	/api/instances/Sds::{id}/action/addSdsRfcacheDevice
--add_standby_mdm	/api/instances/System/action/addStandbyMdm
--add_storage_pool	/api/types/StoragePool/instances
--add_user	/api/types/User/instances
--add_volume	/api/types/Volume/instances
--allow_commands_during_upgrade	/api/instances/System/action/allowCommandsDuringUpgrade
--approve_all_mdm_certificates	/api/getHostCertificate/Mdm?host=<host_ip>
--assign_ldap_groups_to_roles	/api/instances/System/action/assignRoleToLdapGroup
--clear_fault_set	/api/types/FaultSet::{id}/instances/action/clearFaultSet

Table 32 ScaleIO CLI commands (continued)

ScaleIO CLI command	REST API URI
--clear_sds_device_error	/api/instances/Sds::{id}/action/clearDevicesError
--clear_sds_rfcache_error	/api/instances/Sds::{id}/action/clearRfcacheError
--create_mdm_cluster	/api/instances/System/action/createMdmCluster
--delete_user	/api/instances/User::{id}/action/removeUser
--disable_admin	/api/instances/System/action/disableAdmin
--disable_background_device_scanner	/api/instances/StoragePool::{id}/action/disableBackgroundDeviceScanner
--disable_sds_rfcache	/api/instances/ProtectionDomain::{id}/action/disableSdsRfcache
--disable_sds_rmcache	/api/instances/ProtectionDomain::{id}/action/setSdsRmcacheEnabled
	/api/instances/Sds::{id}/action/setSdsRmcacheEnabled
--enable_background_device_scanner	/api/instances/StoragePool::{id}/action/enableBackgroundDeviceScanner
--enable_sds_rfcache	/api/instances/ProtectionDomain::{id}/action/enableSdsRfcache
--enable_sds_rmcache	/api/instances/ProtectionDomain::{id}/action/setSdsRmcacheEnabled
	/api/instances/Sds::{id}/action/setSdsRmcacheEnabled
--enter_maintenance_mode	/api/instances/System/action/enterMaintenanceMode
--exit_maintenance_mode	/api/instances/System/action/exitMaintenanceMode
--finalize_upgrade	/api/instances/System/action/finalizeUpgrade
--generate_certificate	/api/instances/Sds::{id}/action/generateCertificate
--generate_mdm_certificate	/api/instances/System/action/generateMdmCertificate
--generate_mdm_csr_file	/api/instances/System/action/generateMdmCsrFile
--inactivate_protection_domain	/api/instances/ProtectionDomain::{id}/action/inactivateProtectionDomain
--logout	/api/logout
--map_volume_to_sdc	/api/instances/Volume::{id}/action/addMappedSdc
--modify_cluster_virtual_ips	/api/instances/System/action/modifyClusterVirtualIPs
--modify_management_ip	/api/instances/System/action/setManagementIps
--modify_sds_device_capacity	/api/instances/Device::{id}/action/setDeviceCapacityLimit
--modify_sds_ip_role	/api/instances/Sds::{id}/action/setSdsIpRole
--modify_sds_port	/api/instances/Sds::{id}/action/setSdsPort
--modify_spare_policy	/api/instances/StoragePool::{id}/action/setSparePercentage
--modify_user	/api/instances/User::{id}/action/setUserRole
--modify_virtual_ip_interfaces	/api/instances/System/action/modifyVirtualIp

Table 32 ScaleIO CLI commands (continued)

ScaleIO CLI command	REST API URI
--modify_volume_capacity	/api/instances/Volume::{id}/action/setVolumeSize
--modify_zero_padding_policy	/api/instances/StoragePool::{id}/action/setZeroPaddingPolicy
--query_all	/api/instances
--query_all_approved_sdc	/api/types/Sdc/instances
--query_all_device_latency_meters	/api/types/Device/instances/action/querySelectedStatistics
--query_all_fault_sets	/api/types/FaultSet/instances
--query_all_sdc	/api/types/Sdc/instances
--query_all_sds	/api/types/Sds/instances
--query_all_volumes	/api/types/Volume/instances
--query_cluster	/api/instances/System/queryMdmCluster
--query_device_latency_meters	/api/types/Device/instances/action/querySelectedStatistics
--query_device_test	/api/instances/ProtectionDomain::{id}/action/queryDevicesTestResults
--query_fault_set	/api/instances/FaultSet::{id}
--query_license	/api/instances/System::{id}/
--query_network_latency_meters	/api/instances/Sds::{id}/action/querySdsNetworkLatencyMeters
--query_oscillating_failure_counter_parameters failure_counter	/api/instances
--query_performance_parameters	api/instances
--query_properties	/api/instances/querySelectedStatistics
--query_protection_domain	/api/instances/ProtectionDomain::{id}/
--query_remote_read_only_limit_state	/api/instances
--query_remote_syslog	/api/Configuration
--query_restricted_sdc_mode	/api/instances/System::{id}
--query_sdc	/api/instances/Sdc::{id}
	/api/instances/Sdc::{id}/action/queryOscillatingNetworkCounters
	/api/instances/System/action/querySdcOscillatingNetworkCounters
--query_sdc_volume_limits	/api/instances/Volume::{id}
--query_sds	/api/instances/Sds::{id}/
	/api/instances/Sds::{id}/action/queryOscillatingNetworkCounters
	/api/instances/System/action/querySdsOscillatingNetworkCounters
--query_sds_device_info	api/instances/Device::{id}
	/api/instances/Sds::{id}/relationships/Device
--query_sds_network_test_results	/api/instances/Sds::{id}/action/querySdsNetworkTestResults

Table 32 ScaleIO CLI commands (continued)

ScaleIO CLI command	REST API URI
--query_storage_pool	/api/instances/StoragePool::{id}/
--query_system_limits	/api/instances/System/action/querySystemLimits
--query_upgrade	/api/types/System/instances
--query_user	/api/instances/User::{id}
--query_user_authentication_properties	/api/instances/System/action/queryLdap
--query_users	/api/types/User/instances
--query_volume	/api/instances/Volume::{id}
--query_volume_tree	/api/instances/VTree::{id}
--refresh_mdm_cluster_capabilities	/api/instances/System/action/refreshMdmClusterCapabilities
--remove_consistency_group_snapshots	/api/instances/System::{id}/action/ removeConsistencyGroupSnapshots
--remove_fault_set	/api/instances/FaultSet::{id}/action/removeFaultSet
--remove_ldap_group_from_role_assignment	/api/instances/System/action/removeLdapGroup
--remove_ldap_service	/api/instances/System/action/deleteLdapService
--remove_protection_domain	/api/instances/ProtectionDomain::{id}/action/ removeProtectionDomain
--remove_sdc	/api/instances/Sdc::{id}/action/removeSdc
--remove_sds	/api/instances/Sds::{id}/action/removeSds
--remove_sds_device	/api/instances/Device::{id}/action/removeDevice
--remove_sds_ip	/api/instances/Sds::{id}/action/addSdsIp
--remove_sds_rfcache_device	/api/instances/RfcacheDevice::{id}/action/removeSdsRfcacheDevice
--remove_standby_mdm	/api/instances/System/action/removeStandbyMdm
--remove_storage_pool	/api/instances/StoragePool::{}/action/removeStoragePool
--remove_volume	/api/instances/Volume::{id}/action/removeVolume
--rename_device	/api/instances/Device::{id}/action/setDeviceName
--rename_fault_set	/api/types/FaultSet::{id}/instances/action/setFaultSetName
--rename_ldap_service	/api/instances/System/action/renameLdapService
--rename_mdm	/api/instances/System/action/renameMdm
--rename_protection_domain	/api/instances/ProtectionDomain::{id}/action/ setProtectionDomainName
--rename_sdc	/api/instances/Sdc::{id}/action/setSdcName
--rename_sds	/api/instances/Sds::{id}/action/setSdsName
--rename_sds_rfcache_device	/api/instances/RfcacheDevice::{id}/action/renameSdsRfcacheDevice
--rename_storage_pool	/api/instances/StoragePool::{id}/action/setStoragePoolName

Table 32 ScaleIO CLI commands (continued)

ScaleIO CLI command	REST API URI
--rename_system	/api/instances/System::{id}/action/setSystemName
--rename_volume	/api/instances/Volume::{id}/action/setVolumeName
--replace_cluster_mdm	/api/instances/System/action/replaceClusterMdm
--replace_mdm_security_files	/api/instances/System/action/replaceMdmSecurityFiles
--reset_oscillating_failure_counters	/api/instances/System/action/resetSdsOscillatingCounter
	/api/instances/System/action/resetSdcOscillatingCounter
	/api/instances/StoragePool::{id}/action/ resetDeviceOscillatingCounter
	/api/instances/ProtectionDomain::{id}/action/resetOscillatingCounter
--reset_password	/api/instances/User::{id}/action/resetPassword
--reset_scanner_error_counters	/api/instances/StoragePool::{id}/action/ resetBackgroundDeviceScannerErrorCounters
--set_capacity_alerts_threshold	/api/instances/ StoragePool::{id}/action/setCapacityAlertThresholds
--set_checksum_mode	/api/instances/StoragePool::{id}/action/setChecksumEnabled
--set_component_authentication_properties	/api/instances/System/action/setConnectionSecurityPolicy
--set_drl_properties	/api/instances/System::{id}/action/setDrlMode
	api/instances/Sds::{id}/action/setDrlMode
--set_license	/api/instances/System::{id}/action/setLicense
--set_login_banner	/im/types/Configuration/actions/setLoginBanner
--set_management_client_communication	/api/instances/System/action/ enableManagementClientSecureCommunication
--set_obfuscation_properties	/api/instances/System::{id}/action/setObfuscationProperties
--set_oscillating_failure_counter_parameters	/api/instances/System/action/setSdsOscillatingCounterParameters
	/api/instances/System/action/setSdcOscillatingCounterParameters
	/api/instances/ProtectionDomain::{id}/action/ setOscillatingCounterParameters
	/api/instances/System/action/ setDeviceOscillatingCounterParameters
--set_password	/api/instances/User/action/setPassword
--set_performance_parameters	/api/instances/System/action/setMdmPerformanceParameters
	/api/instances/ProtectionDomain::{id}/action/ setSdsPerformanceParameters
	/api/instances/Sds::{id}/action/setSdsPerformanceParameters
	/api/instances/Sdc::{id}/action/setSdcPerformanceParameters
	/api/instances/System/action/setSdsPerformanceParameters

Table 32 ScaleIO CLI commands (continued)

ScaleIO CLI command	REST API URI
	/api/instances/System/action/setSdcPerformanceParameters
	/api/instances/FaultSet::{id}/action/setSdsPerformanceParameters
--set_rebalance_mode	/api/instances/StoragePool::{id}/action/setRebalanceEnabled
--set_rebalance_policy	/api/instances/StoragePool::{id}/action/setRebalanceIoPriorityPolicy
--set_rebuild_mode	/api/instances/StoragePool::{id}/action/setRebuildEnabled
--set_rebuild_policy	/api/instances/StoragePool::{id}/action/setRebuildIoPriorityPolicy
--set_rebuild_rebalance_parallelism	/api/instances/StoragePool::{id}/action/ setRebuildRebalanceParallelism
--set_restricted_sdc_mode	/api/instances/System::{id}/action/setRestrictedSdcModeEnabled
--set_rfcache_parameters	/api/instances/ProtectionDomain::{id} /action/setRfcacheParameters
--set_rfcache_usage	/api/instances/StoragePool::{id}/action/setRfcacheUsage
--set_rmcache_usage	/api/instances/StoragePool::{id}/action/setUseRmcache
--set_rmcache_write_handling_mode	/api/instances/StoragePool::{id}/action/ setRmcacheWriteHandlingMode
--set_sdc_volume_limits	/api/instances/Volume::{id}/action/setMappedSdcLimits
--set_sds_network_limits	/api/instances/ProtectionDomain::{id}/action/setSdsNetworkLimits
--set_sds_rmcache_size	/api/instances/ProtectionDomain::{id}/action/setSdsRmcacheSize
	/api/instances/Sds::{id}/action/setSdsRmcacheSize
--set_syslog_facility	/api/instances/System::{id}/action/setSyslogFacility
--set_user_authentication_method	/api/instances/System/action/setAuthenticationMethod
--set_volume_rmcache_usage	/api/instances/Volume::{id}/action/setVolumeUseRmcache
--show_certificate	/api/instances/Sds::{id}
--snapshot_volume	/api/instances/System::{id}/action/snapshotVolumes
--start_device_test	/api/instances/ProtectionDomain::{id}/action/startDevicesTest
--start_remote_syslog	/api/instances/System::{id}/action/startRemoteSyslog
--start_sds_network_test	/api/instances/Sds::{id}/action/startSdsNetworkTest
--start_upgrade	/api/instances/System/action/startUpgrade
--stop_remote_syslog	/api/instances/System::{id}/action/stopRemoteSyslog
--switch_cluster_mode	/api/instances/System/action/switchClusterMode
--switch_mdm_ownership	/api/instances/System/action/changeMdmOwnership
--unmap_volume_from_sdc	/api/instances/Volume::{id}/action/removeMappedSdc
--update_device_original_path	/api/instances/Device::{id}/action/updateDeviceOriginalPathname
--update_sds_rfcache_device_original_path	/api/instances/RfcacheDevice::{id}/action/ updateSdsRfcacheDeviceOriginalPath

Software type resources

This section contains information about software type resources.

General information

This section describes general information for the ScaleIO REST API, including REST API version in header fields, URIs, HTTP authentication, response fields, and ScaleIO Gateway configuration.

REST API version

The REST API version should be added to the `Accept` parameter in the header field. For example: `Accept: application/json;version=1.0`

If no Accept header field is present, it is assumed that the client accepts all media types and versions.

HTTP authentication

To perform authentication, invoke the HTTP GET request, with the URI: `/api/login` , with `MDM user\password` passes in HTTP Basic authentication. A token is returned.

For example:

```
https://10.76.60.190:443/api/login
```

RESPONSE

200 OK

HEADERS form

Cache-Control:	no-cache
Cache-Control:	no-store
Content-Encoding:	gzip
Content-Type:	application/json;charset=UTF-8
Date:	2014 Aug 21 14:18:38
Expires:	1970 Jan 1 02:00:00 -44 years
Pragma:	no-cache
Server:	Apache-Coyote/1.1
Transfer-Encoding:	chunked
Vary:	Accept-Encoding

BODY

"YWRtaW46MTQwODY0ODcxODYwMTpmYzRmZDg1MTFmYmNiOGU5YzQ5ZTFmOGZjNTIxOTcxYw"

download

This token is valid for 8 hours from the time it was created, unless there has been no activity for 10 minutes, or if the client has sent a logout request.

HTTP token invalidation (logout)

The invalidation is done by invoking the HTTP GET request, with the URI: `/api/logout`

The token mentioned above is the password in the HTTP Basic authentication (the user is ignored - it can be empty).

For every REST API request that does not require the `gatewayAdminPassword`, the authentication is done by passing the token mentioned above as the password in the HTTP Basic authentication (the user is ignored - it can be empty).

Requests that require the `gatewayAdminPassword` work similarly, except that instead of `/api/login`, invoke an HTTP GET request, `/api/gatewayLogin` with `user: admin password: <gatewayAdminPassword>` in HTTP Basic authentication. A token is returned. Instead of invoking `/api/logout`, invoke `/api/gatewayLogout` with the token received when you logged in.

Note

Requests that require `gatewayAdminPassword` are:
GET:

```
/api/Configuration
/api/gatewayLogout
/api/getHostCertificate/{Mdm | Lia}
```

POST:

```
/api/updateConfiguration
/api/instances/System/action/createMdmCluster
/api/trustHostCertificate/{Mdm | Lia}
/api/gatewaySetSecureCommunication
```

Response fields

The order of the fields in the responses may change. More fields may be added in the future.

URIs

- POST (create) / GET all objects for a given type:

  `/api/types/{type}/instances`
- GET by id:

  `/api/instances/{type::id}`
- POST a special action on an object:

  `/api/instances/{type::id}/action/{actionName}`
- POST a special action on a given type:

  `/api/types/{type}/instances/action/{actionName}`
- Get current API version:

  `/api/version`
- Every row in the Object's Parent table appears as a link in the response of get object:

  `/api/instances/{type::id}`
- Every row in the Object's Relationships table appears as a link in the response of get object:

  `/api/instances/{type::id}/relationships/{Relationship name}`
- GET all instances

  `/api/instances/`

Table 33 Response

Property	Type	Note
sessionTag	Long	

Table 33 Response (continued)

Property	Type	Note
isDirty	Boolean	The version on some objects changed while the MDM created the response
System	Syst	
lastSystemVersion	Long	
protectionDomainList	List of Protection Domain objects	
lastProtectionDomainVersion	Long	
sdsList	List of SDS objects	
lastSdsVersion	Long	
storagePoolList	List of StoragePool objects	
lastStoragePoolVersion	Long	
deviceList	List of Device objects	
lastDeviceVersion	Long	
volumeList	List of Volume objects	
lastVolumeVersion	Long	
vTreeList	List of VTree objects	
lastVTreeVersion	Long	
sdcList	List of SDC objects	
lastSdcVersion	Long	
faultSetList	List of Fault Set objects	
lastFaultSetVersion	Long	
rfcacheDeviceList	List of RfcacheDevice objects	
rfcacheDeviceVersion	Long	

- Change configuration of ScaleIO Gateway (POST)

  `/api/updateConfiguration/`

  Request:

  Parameters (AND\OR):

  - `mdmAddresses`—List of MDM IP addresses that the ScaleIO Gateway will use to connect to the MDM (for performance improvement, place the master MDM IP addresses before the slave ones)

- `mdmPort`—MDM port
- `gatewayAdminPassword`—Password for installation manager, get/update ScaleIO gateway configuration, and add master MDM
- `systemId`—ID of the system configured by `mdmAddresses`. If `systemId` is left empty, it will be populated on the first invocation of `/api/login`. At every /api/login, the system ID value is compared to the ID of the ScaleIO cluster (MDM). If the IDs are different (`mdmAddresses` points to a cluster that does not match the `systemId`), an error message is displayed, and all login tokens (/api/login, /api/gatewayLogin) are invalidated. None of the gateway clients will be able to send requests to the MDM until one of the following happens:
  - `systemId` is configured to the correct value
  - `mdmAddresses` is configured to the correct IP addresses
  - `systemId` is set to an empty value. In this case, on the first invocation of `/api/login`, the value for `systemId` will be populated according to the ID of the matching `mdmAddresses`.

    If the value for `systemId` is different from the configured `systemId` (and the configured `systemId` is not empty), all login tokens (/api/login, /api/gatewayLogin) are invalidated.
- `snmpSamplingFrequency`—MDM sampling frequency in seconds. If sampling frequency is set to `0`, the SNMP trap sender will be disabled, and no sampling to the MDM will be performed.
- `snmpResendFrequency`—Resend frequency of SNMP traps, in minutes. If resend frequency is set to `0`, all traps will be sent in every sampling.
- `snmpTrapsReceiverIp`—SNMP trap receivers' IP addresses (supports up to two comma-separated or semi-colon-separated IP addresses)
- `snmpPort`—The port number used for SNMP traps

For example:

```
{
"mdmAddresses":["10.76.60.150", "10.76.60.11"],
"mdmPort":"6611",
"gatewayAdminPassword":"Password1",
"systemId":"7f5d8fc72a3d7f3d" ,
"snmpSamplingFrequency":"30",
"snmpResendFrequency":"0",
"snmpTrapsReceiverIps":["10.76.60.190","10.76.60.191"],
"snmpPort":"162"
}
```

- Get configuration of ScaleIO Gateway (GET)

  `/api/Configuration/`

  Response:

  - `mdmAddresses`—List of MDM master IP addresses that the ScaleIO Gateway will use to connect to the MDM
  - `mdmPort`—MDM port
  - `mdmUsername`—MDM user name

- `systemId`—ID of the system, configured by `mdmAddresses` (see detailed explanation above, for "Change configuration of ScaleIO Gateway (POST)".)
- `snmpSamplingFrequency`—MDM sampling frequency in seconds
- `snmpResendFrequency`—Resend frequency of SNMP traps, in minutes
- `snmpTrapsReceiverIp`—SNMP trap receivers' IP addresses
- `snmpPort`—The port number used for SNMP traps
- `remoteSyslog`—The remote syslog servers configuration in the MDM
- `featuresEnableSnmp`—Indicates whether SNMP is enabled or not
- `cipherSuites`—A list of names of cipher suites (as recognized by Java) that the SSL will use instead of the default list.
- `featuresEnableIM`—Indicates whether the Installation Manager is enabled or not
- `allowNonSecureCommunication`

  For example:

```
{
"snmpResendFrequency": "0",
"snmpSamplingFrequency": "30",
"snmpPort": "162",
"mdmPort":"6611",
"remoteSyslog": [
       {
           "hostName": "10.76.60.100",
           "port": 1468,
           "facility": 16
       },
       {
           "hostName": "10.76.60.101",
           "port": 1468,
           "facility": 16
       }
],
"mdmAddresses":[
"10.76.60.150",
"10.76.60.135"
],
"systemId":"7f5d8fc72a3d7f3d",
"snmpTrapsReceiverIps": [
       "10.76.60.192"
   ]
}
```

- Query selected statistics (POST)

  `/api/instances/querySelectedStatistics`

  Request:

  - Required parameters:

    `selectedStatisticsList`—list of objects containing:

    - `type`—object type (System, ProtectionDomain, Sds, StoragePool, Device, Volume, VTree, Sdc, FaultSet, RfcacheDevice)
    - `properties`—list of properties to fetch
  - The following parameters are not relevant to the System type (can be omitted):

`ids`—list of objects ids

or

`allIds`—with empty value

Response:

The response contains a list of all requested types, the ids, and properties, according to the request parameters.

The System type does not contain an id.

Types for which the requested ids do not exist in the cluster will appear with an empty {} - see `FaultSet` in the example below.

Any id (from the ids parameters in the request) that does not exist in the cluster will not appear in the response. The following will appear in the scaleio.log:

`Got no statistics for <type>::<id>.`

In the following example, `Sds::022beb2300000004` does not exist in the cluster, so it does not appear in the response. Instead, this will appear in the scaleio.log:

```
Got no statistics for Sds::id: 022beb2300000004
```

Example:

`POST https://localhost:8443/api/instances/querySelectedStatistics`

Body:

```
{"selectedStatisticsList":[
{"type":"ProtectionDomain", "ids":["cc480c9b00000000"],
"properties":["capacityInUseInKb">,
 {"type":"Volume", "ids":
["022beb2500000006","022beb2300000004"], "properties":
["numOfMappedSdcs", "userDataWriteBwc">,
{"type":"Sds", "ids":["c919d82000000001","022beb2300000004"],
"properties":["capacityInUseInKb">,
 {"type":"System", "allIds":"", "properties":
["rmcacheSizeInKb">,
 {"type":"FaultSet", "ids":
["c919d82000000001","022beb2300000004"], "properties":
["numOfSds">,
{"type":"StoragePool", "allIds":"", "properties":
["unreachableUnusedCapacityInKb", "numOfThinBaseVolumes">
 >
```

The response:

```
{
        "FaultSet":{
        },
        "Volume":{
                "022beb2300000004":{
                        "userDataWriteBwc":{
                                "numOccured":0,
                                "totalWeightInKb":0,
                                "numSeconds":1
                        },
                        "numOfMappedSdcs":0

                },
                "022beb2500000006":{
                        "userDataWriteBwc":{
                                "numOccured":0,
                                "totalWeightInKb":0,
                                "numSeconds":1
                        },
                        "numOfMappedSdcs":0

                }
        },
        "StoragePool":{
                "7db925c200000003":{
                        "numOfThinBaseVolumes":0,
                        "unreachableUnusedCapacityInKb":0
                },
                "7db925bf00000000":{
                        "numOfThinBaseVolumes":2,
                        "unreachableUnusedCapacityInKb":0
                },
                "7db925c000000001":{
                        "numOfThinBaseVolumes":0,
                        "unreachableUnusedCapacityInKb":0
                },
                "7db925c100000002":{
                        "numOfThinBaseVolumes":0,
                        "unreachableUnusedCapacityInKb":0
                }
        },
        "ProtectionDomain":{
                "cc480c9b00000000":{
                        "capacityInUseInKb":50331648
                }
        },
        "System":{
                "rmcacheSizeInKb":524288
        },
        "Sds":{
                "c919d82000000001":{
                        "capacityInUseInKb":12509184
                }
        }
}
```


- Get MDM's or LIA's certificate (GET)

```
/api/getHostCertificate/{Mdm|Lia}?host={host ip}
```

The request can be sent for either MDM or LIA.

Example:

```
/api/getHostCertificate/Mdm?host=10.76.60.10
```

Response:

The host certificate in PEM encoding.

Note

The whole certificate should be saved into a `cer` file for the `trustHostCertificate`request.

For example:


```
-----BEGIN CERTIFICATE-----
MIIDcTCCAlmgAwIBAAIBATANBgkqhkiG9w0BAQUFADB8MQwwCgYDVQQqEwNNRE0xFzA
VBgNVBAMTDmNlbnRvcy02LTQtYWRpMRIwEAYDVQQHEwlIb3BraW50b24xFjAUBgNVBA
gTDU1hc3NhY2h1c2V0dHMxCzAJBgNVBAYTAlVTMQwwCgYDVQQKEwNFTUMxDDAKBgNVB
AsTA0FTRDAeFw0xNTEyMTkwNzI4MTVaFw0yNTEyMTcwODI4MTVaMHwxDDAKBgNVBCoT
A01ETTEXMBUGA1UEAxMOY2VudG9zLTYtNC1hZGkxEjAQBgNVBAcTCUhvcGtpbnRvbjE
WMBQGA1UECBMNTWFzc2FjaHVzZXR0czELMAkGA1UEBhMCVVMxDDAKBgNVBAoTA0VNQz
EMMAoGA1UECxMDQVNEMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEA4SMyb
aAEZjfBX9wLglr3wxYHOvID5Pe1Z26Pv8oIR/
MTOVa1Bw4A9px1MHHSIfkAfgRlLC24uebZXhbh0snBq+OL+SJPwEfbOVbif/
saXL8RJFwm/VNg8KHUwjuq/sJkKDjx9uSf0U+/9FzwvKVuM87xDj/
rVvJgBYh6pH34q/XD5l8am/iEQr/EnGZmIsa+VkCL0IeYKbkA3ZINfI4YsjSJ
+qeu5e/
KMsNlHEvmhk1DdJbLayn9QkiS5Q9e8A40jjkb2e1Q71awoOlb6+8XXWWkpBhxAnRa9P
8Pb1BfcNyUfXtrKuy+fRjw4Gp
+rw2MdoIDuMbO1+1sQaRvVPTYxwIDAQABMA0GCSqGSIb3DQEBBQUAA4IBAQCu1e/
jBz63ZFDS+pFvZ3XI/VMdn9NArM
+8Cjn7Luar8oEVAYI6jqYYcZCk2jQyfuI1HP2jXqPTJo8CNyNT5S6rZ5ryHOiRjn/
K2pVC6kT497LY5lc3LjhXUdLjpWnW2jsGfM93cCkkrxu8wmkh9oo8WizOiRAyKmz02u
TEuEok7GJBS/
DR6csnLo2YLUV6ZqeBN9jdzZbIY7SoFWya1K4xZmqhkAtnj1ynP3uoxTkd
+wfDRmYeDv8l5eciLj2BXNuV8zXYWSCyABZC//jvajNtSEXgUura3uh0YBIfbO/
AZ980zUMwJBMBr06yw4tHnHRRYgfI3tnZOD4byaJOdHuq
-----END CERTIFICATE-----
```


- Trust an MDM's or LIA's certificate (POST)

```
/api/trustHostCertificate/{Mdm|Lia}
```

Example:

```
/api/trustHostCertificate/Mdm
```

Request:

Content-type: multipart/form-data

The body should contain a part named "file" and the file containing the certificate to be trusted.

- Set the gateway to work with secured communication with the MDM (POST):

```
/api/gatewaySetSecureCommunication
```

  The gateway will not be able to connect to the MDM using non-secured communication.

- Working with the Installation Manager (IM) REST API

  IM is an orchestration engine that runs commands from different queues (nodes and MDM) at different "phases," and labels them as a "process" (such as an upgrade).

  The IM REST API begins with the `/im/`prefix.

  Login to the IM REST API must be done via POST to the j_spring_security_check url, followed by username, password, and the login submission. For example:

```
POST https://localhost/j_spring_security_check
Content: "j_username=admin&j_password=Password1&submit=Login"
```

  An operation such as the above will provide a JSESSIONID cookie in the response that should be used for identification of the session in all the future requests.

  For example:

```
POST "Content-Type: application/json" "Cookie:
JSESSIONID=969F624A761937AE80E6CC9E91756B10" https://
localhost/im/types/Command/instances/actions/retry
```

  - Most IM operations require a topology / configuration object as a parameter (payload) for the REST (HTTP) request in JSON format.

    Such an object has the following format:

    Note

    Each field's functionality is marked with the /* */ notation—it should not be used when generating such instance with JSON format `{`

    /* target IO address to send SNMP traps while configuring Gateway to become SNMP sampler */

```
    "snmpIp": null,
```

    /* list of IP addresses from which LIA will only accept connections—if null, accept any connection */

```
    "safeIPsForLia": null,
```

/* system installation ID to be used or configured*/

```
    "installationId": null,
```

/* all MDM manager IP addresses list*/

```
    "mdmIPs": ["10.76.60.48",
    "10.76.60.41"],
```

/* password for MDM and LIA */

```
    "mdmPassword": "Password1",
    "liaPassword": "Password1",
```

/* deprecated—not in use */

```
    "licenseKey": null,
```

/* deprecated—not in use */

```
    "licenseType": "unlimitedCapacity",
```

/* boolean, indicating whether cluster should be set with high performance profile */

```
    "isClusterOptimized": null,
```

/* deprecated—not in use */

```
    "callHomeConfiguration": null,
```

/* string describing parameters for remote sys log configuration */

```
    "remoteSyslogConfiguration": null,
```

/* string describing cluster (master) version */

```
    "systemVersionName": "",
```

/* boolean indicating whether system is running */

```
    "upgradeRunning": false,
```

/* boolean indicating whether current MDM is a cluster (if running) */

```
    "clustered": false,
    "masterMdm": {
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.41"],
```

/* domain only relevant (if at all) for Windows */

```
            "domain": null,
            "userName": "root",
            "password": "Password1",
            "liaPassword": null
        },
```

/* internal data fields—do not use */

```
        "nodeInfo": null,
```

/* in case of upgrade—if rollback should occur—to which version */

```
        "rollbackVersion": null,
        "mdmIPs": ["10.76.60.41"],
        "name": "M2",
```

/* unique MDM ID */

```
        "id": null,
```

/* internal data fields—do not use */

```
        "ipForActor": null,
```

/* IP addresses to be used and configured for management */

```
        "managementIPs": ["10.76.60.41"]
    },
```

/* similar to master MDM */

```
    "slaveMdmSet": [{
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.48"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
```

```
            "liaPassword": null
        },
        "nodeInfo": null,
        "rollbackVersion": null,
        "mdmIPs": ["10.76.60.48"],
        "name": "M1",
        "id": null,
        "ipForActor": null,
        "managementIPs": ["10.76.60.48"]
    }],
```

/* similar to master MDM */

```
    "tbSet": [{
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.146"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
            "liaPassword": null
        },
        "nodeInfo": null,
        "rollbackVersion": null,
        "mdmIPs": ["10.76.60.146"],
        "name": "M3",
        "id": null,
        "tbIPs": ["10.76.60.146"]
    }],
```

/* similar to master MDM */

```
    "standbyMdmSet": [],
```

/* similar to master MDM */

```
    "standbyTbSet": [],
    "sdsList": [{
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.146"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
            "liaPassword": null
        },
        "nodeInfo": null,
        "rollbackVersion": null,
        "sdsName": "SDS_10.76.60.146",
```

/* name (required) and id (optional, if it already exists) for Protection Domain */

```
        "protectionDomain": "domain1",
        "protectionDomainId": null,
```

/* fault set id (if defined) */

```
        "faultSet": null,
```

/* various SDS IP addresses and their usage */

```
        "allIPs": ["10.76.60.146"],
        "sdsOnlyIPs": null,
        "sdcOnlyIPs": null,
        "devices": [],
```

/* Read Flash Cache definition if defined, on which pools and which devices */

```
        "rfCached": false,
        "rfCachedPools": [],
        "rfCachedDevices": [],
```

/* boolean, indicating whether SDS is set to high performance profile */

```
        "optimized": false,
        "port": 7072,
        "id": "0"
    },
    {
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.48"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
            "liaPassword": null
        },
        "nodeInfo": null,
        "rollbackVersion": null,
        "sdsName": "SDS_10.76.60.48",
        "protectionDomain": "domain1",
        "protectionDomainId": null,
        "faultSet": null,
        "allIPs": ["10.76.60.48"],
        "sdsOnlyIPs": null,
        "sdcOnlyIPs": null,
        "devices": [],
        "rfCached": false,
        "rfCachedPools": [],
        "rfCachedDevices": [],
        "optimized": false,
        "port": 7072,
        "id": "0"
    },
    {
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.41"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
```

```
            "liaPassword": null
        },
        "nodeInfo": null,
        "rollbackVersion": null,
        "sdsName": "SDS_10.76.60.41",
        "protectionDomain": "domain1",
        "protectionDomainId": null,
        "faultSet": null,
        "allIPs": ["10.76.60.41"],
        "sdsOnlyIPs": null,
        "sdcOnlyIPs": null,
        "devices": [],
        "rfCached": false,
        "rfCachedPools": [],
        "rfCachedDevices": [],
        "optimized": false,
        "port": 7072,
        "id": "0"
    }],
    "sdcList": [{
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.41"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
            "liaPassword": null
        },
        "nodeInfo": null,
        "rollbackVersion": null,
```

/* splitter (RecoverPoint) IP address to replicate the IO */

```
"splitterRpaIp": null,
```

/* Boolean indicating whether SDC should be set to high performance profile */

```
        "optimized": false
    },
    {
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.48"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
            "liaPassword": null
        },
        "nodeInfo": null,
        "rollbackVersion": null,
        "splitterRpaIp": null,
        "optimized": false
    }],
```

/* security setting—whether to allow process to connect to non-secured LIA / MDM, and whether ScaleIO components should be configured to be non-secured (secured by default) */

```
    "securityConfiguration": {
        "allowNonSecureCommunicationWithMdm": false,
        "allowNonSecureCommunicationWithLia": false,
        "disableNonMgmtComponentsAuth": false
    },
}
```

- The Node object is used as an attribute of each ScaleIO component (for example, on which node the MDM is installed), and describes how the node will be accessed:
  - via SSH (Linux node with root credentials)
  - via WMI (Windows node with admin domain and credentials)
  - via LIA (unknown node with null domain, null credentials and LIA password)

    For example:

```
//linux node accessed via SSH
"node": {
            "ostype": "unknown",
            "nodeName": null,
            "nodeIPs": ["10.76.60.48"],
            "domain": null,
            "userName": null,
            "password": null,
            "liaPassword": null
        },
//unknown node (linux or windows) - accessed via LIA
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.41"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
            "liaPassword": null
        },
```

- LIA usage—Any post installation operations MUST be done via LIA. Such operations (for example, upgrades) may fail unless all nodes are targeted to be performed via LIA.

System

Table 34 System Objects

Properties	Type
id	String
name	String
systemVersionName	String
primaryMdmActorIpList (deprecated)	List
primaryMdmActorPort (deprecated)	Integer

Table 34 System Objects (continued)

Properties	Type
secondaryMdmActorIpList (deprecated)	List
secondaryMdmActorPort (deprecated)	Integer
tiebreakerMdmIpList (deprecated)	List
tiebreakerMdmPort (deprecated)	Integer
mdmMode (deprecated)	Single or Cluster
mdmClusterState (deprecated)	NotClustered or ClusteredNormal or ClusteredDegraded or ClusteredTiebreakerDown or ClusteredDegradedTiebreakerDown
mdmManagementIpList (deprecated)	List
mdmManagementPort	Integer
capacityAlertHighThresholdPercent	Default value for newly created Storage Pool.
capacityAlertCriticalThresholdPercent	Default value for newly created Storage Pool.
installId	String
swid	String
daysInstalled	Long
maxCapacityInGb	String
capacityTimeLeftInDays	String
enterpriseFeaturesEnabled	Boolean
defaultIsVolumeObfuscated	Deprecated in v2.0
isInitialLicense	Boolean
mdmCluster	MdmCluster
upgradeState	NoUpgrade, MdmUpgradeInProgress or UpgradeInProgress
performanceParameters	MdmPerfParams
currentProfilePerformanceParameters	MdmPerfParams
cliPasswordAllowed	Boolean
managementClientSecureCommunicationEnabled	Boolean
tlsVersion	String
showGuid	Boolean

Table 34 System Objects (continued)

Properties	Type
authenticationMethod	Native, Ldap or NativeAndLdap
mdmToSdsPolicy	None or Authentication
sdcMdmNetworkDisconnectionsCounterParameters	OscillatingCounterParameters
sdcSdsNetworkDisconnectionsCounterParameters	OscillatingCounterParameters
sdcMemoryAllocationFailuresCounterParameters	OscillatingCounterParameters
sdcSocketAllocationFailuresCounterParameter	OscillatingCounterParameters
sdcLongOperationsCounterParameters	OscillatingCounterParameters
restrictedSdcModeEnabled	Boolean
remoteReadOnlyLimitState	Boolean

Table 35 System Relationships

Name	Type	Note
ProtectionDomain	List of Protection Domain objects	
SDC	List of SDC objects	
User	List of User objects	
Statistics		

MDM Cluster

Table 36 MDM Cluster Object

Properties	Type
id	String
name	String
clusterMode	Invalid, OneNode, ThreeNodes or FiveNodes
master	MdmClusterMember
slaves	List<MdmClusterMember>
tieBreakers	List<MdmClusterMember>
standbyMDMs	List<MdmClusterMember>
clusterState	Invalid, ClusteredNormal or ClusteredDegraded

Table 36 MDM Cluster Object (continued)

Properties	Type
goodNodesNum	Integer
goodReplicasNum	Integer

Table 37 MdmClusterMember

Properties	Type
id	String
name	String
port	Integer
ips	List<String>
managementIPs	List<String>
role	Invalid, Manager or TieBreaker
status	Invalid, Normal, Error, Disconnected or OutOfSync
versionInfo	String

Table 38 MdmPerfParams

Properties	Type
mdmNumberSdcReceiveUmt	Integer
mdmNumberSdsReceiveUmt	Integer
mdmNumberSdsSendUmt	Integer
mdmNumberSdsKeepaliveReceiveUmt	Integer
mdmSdsCapacityCountersUpdateInterval	Integer
mdmSdsCapacityCountersPollingInterval	Integer
mdmSdsVolumeSizePollingInterval	Integer
mdmSdsVolumeSizePollingRetryInterval	Integer
mdmNumberSdsTasksUmt	Integer
mdmInitialSdsSnapshotCapacity	Integer
mdmSdsSnapshotCapacityChunkSize	Integer
mdmSdsSnapshotUsedCapacityThreshold	Integer
mdmSdsSnapshotFreeCapacityThreshold	Integer
perfProfile	Default or HighPerformance

Table 39 OscillatingCounterParameters

Properties	Type
shortWindow	OscillatingCounterWindowParameters
mediumWindow	OscillatingCounterWindowParameters
longWindow	OscillatingCounterWindowParameters

Table 40 OscillatingCounterWindowParameters

Properties	Type
threshold	Integer
windowSizeInSec	Integer

System Statistics and Operations

Table 41 System Statistics Object

Properties	Type	Note
numOfProtectionDomains	Long	
numOfSdc	Long	
numOfSds	Long	
numOfVolumes	Long	
numOfDevices	Long	
numOfVtrees	Long	
numOfStoragePools	Long	
rmcacheSizeInKb	Long	
capacityLimitInKb	Long	
maxCapacityInKb	Long	
capacityInUseInKb	Long	
thickCapacityInUseInKb	Long	
thinCapacityInUseInKb	Long	
snapCapacityInUseInKb	Long	
snapCapacityInUseOccupiedInKb	Long	
unreachableUnusedCapacityInKb	Long	
protectedVacInKb	Long	
degradedHealthyVacInKb	Long	

Table 41 System Statistics Object (continued)

Properties	Type	Note
degradedFailedVacInKb	Long	
failedVacInKb	Long	
inUseVacInKb	Long	
activeMovingInFwdRebuildJobs	Long	
pendingMovingInFwdRebuildJobs	Long	
activeMovingOutFwdRebuildJobs	Long	
pendingMovingOutFwdRebuildJobs	Long	
activeMovingInBckRebuildJobs	Long	
pendingMovingInBckRebuildJobs	Long	
activeMovingOutBckRebuildJobs	Long	
pendingMovingOutBckRebuildJobs	Long	
activeMovingInRebalanceJobs	Long	
pendingMovingInRebalanceJobs	Long	
activeMovingRebalanceJobs	Long	
pendingMovingRebalanceJobs	Long	
primaryVacInKb	Long	“primary” represents Master MDMs in versions v2.0 and higher.
secondaryVacInKb	Long	“secondary” represents Slave MDMs in versions v2.0 and higher.
primaryReadBwc	Bandwidth Counter	“primary” represents Master MDMs in versions v2.0 and higher.
primaryReadFromDevBwc	Bandwidth Counter	“primary” represents Master MDMs in versions v2.0 and higher.
primaryWriteBwc	Bandwidth Counter	“primary” represents Master MDMs in versions v2.0 and higher.
secondaryReadBwc	Bandwidth Counter	“secondary” represents Slave MDMs in versions v2.0 and higher.
secondaryReadFromDevBwc	Bandwidth Counter	“secondary” represents Slave MDMs in versions v2.0 and higher.
secondaryWriteBwc	Bandwidth Counter	“secondary” represents Slave MDMs in versions v2.0 and higher.
totalReadBwc	Bandwidth Counter	

Table 41 System Statistics Object (continued)

Properties	Type	Note
totalWriteBwc	Bandwidth Counter	
fwdRebuildReadBwc	Bandwidth Counter	
fwdRebuildWriteBwc	Bandwidth Counter	
bckRebuildReadBwc	Bandwidth Counter	
bckRebuildWriteBwc	Bandwidth Counter	
rebalanceReadBwc	Bandwidth Counter	
rebalanceWriteBwc	Bandwidth Counter	
spareCapacityInKb	Long	
capacityAvailableForVolumeAllocationInKb	Long	
protectedCapacityInKb	Long	
degradedHealthyCapacityInKb	Long	
degradedFailedCapacityInKb	Long	
failedCapacityInKb	Long	
movingCapacityInKb	Long	
activeMovingCapacityInKb	Long	
pendingMovingCapacityInKb	Long	
fwdRebuildCapacityInKb	Long	
activeFwdRebuildCapacityInKb	Long	
pendingFwdRebuildCapacityInKb	Long	
bckRebuildCapacityInKb	Long	
activeBckRebuildCapacityInKb	Long	
pendingBckRebuildCapacityInKb	Long	
rebalanceCapacityInKb	Long	
activeRebalanceCapacityInKb	Long	
pendingRebalanceCapacityInKb	Long	

Table 41 System Statistics Object (continued)

Properties	Type	Note
atRestCapacityInKb	Long	
numOfUnmappedVolumes	Long	
numOfMappedToAllVolumes	Long	
numOfThickBaseVolumes	Long	
numOfThinBaseVolumes	Long	
numOfSnapshots	Long	
numOfVolumesInDeletion	Long	
fixedReadErrorCount	Long	
BackgroundScanCompareCount	Long	
BackgroundScannedInMB	Long	
thinCapacityAllocatedInKm	Long	
rmPendingAllocatedInKb	Long	
semiProtectedVacInKb	Long	
inMaintenanceVacInKb	Long	
activeMovingInNormRebuildJobs	Long	
activeMovingOutNormRebuildJobs	Long	
pendingMovingInNormRebuildJobs	Long	
pendingMovingOutNormrebuildJobs	Long	
primaryReadFromRmcacheBwc	Bandwidth Counter	“primary” represents Master MDMs in versions v2.0 and higher.
secondaryReadFromRmcacheBwc	Bandwidth Counter	“secondary” represents Slave MDMs in versions v2.0 and higher.
normRebuildReadBwc	Bandwidth Counter	
normRebuildWriteBwc	Bandwidth Counter	
rfcacheReadsReceived	Long	
rfcacheWritesReceived	Long	
rfcacheAvgReadTime	Long	
rfcacheAvgWriteTime	Long	
rfcacheSourceDeviceReads	Long	

Table 41 System Statistics Object (continued)

Properties	Type	Note
rfcacheSourceDeviceWrites	Long	
rfacheReadHit	Long	
rfcacheReadMiss	Long	
rfcacheWriteMiss	Long	
rfcacheIosSkipped	Long	
rfcacheReadsSkipped	Long	
rfcacheReadsSkippedAlignedSizeTooLarge	Long	
rfcacheReadsSkippedMaxIoSize	Long	
rfcacheReadsSkippedHeavyLoad	Long	
rfcacheReadsSkippedStuckIo	Long	
rfcacheReadsSkippedLowResources	Long	
rfcacheReadsSkippedInternalError	Long	
rfcacheReadsSkippedLockIos	Long	
rfcacheWritesSkippedMaxIoSize	Long	
rfcacheWritesSkippedHeavyLoad	Long	
rfcacheWritesSkippedStuckIo	Long	
rfcacheWritesSkippedLowResources	Long	
rfcacheWritesSkippedInternalError	Long	
rfcacheWritesSkippedCacheMiss	Long	
rfcacheSkippedUnlinedWrite	Long	
rfcacheIoErrors	Long	
rfcacheReadsFromCache	Long	
rfcacheIosOutstanding	Long	
rfcacheReadsPending	Long	
rfcacheWritePending	Long	
semiProtectedCapacityInKb	Long	
inMaintenanceCapacityInKb	Long	
normRebuildCapacityInKb	Long	
activeNormRebuildCapacityInKb	Long	

Table 41 System Statistics Object (continued)

Properties	Type	Note
pendingNormRebuildCapacityInKb	Long	
numOfRfcacheDevices	Long	
rfcacheFdReadsReceived	Long	
rfcacheFdWritesReceived	Long	
rfcacheFdInlightReads	Long	
rfcacheFdInlightWrites	Long	
rfcacheFdReadTimeGreater500Millis	Long	
rfcacheFdReadTimeGreater1Sec	Long	
rfcacheFdReadTimeGreater5Sec	Long	
rfcacheFdReadTimeGreater1Min	Long	
rfcacheFdWriteTimeGreater500Millis	Long	
rfcacheFdWriteTimeGreater1Sec	Long	
rfcacheFdWriteTimeGreater5Sec	Long	
rfcacheFdWriteTimeGreater1Min	Long	
rfcacheFdAvgReadTime	Long	
rfcacheFdAvgWriteTime	Long	
rfcacheFdIoErrors	Long	
rfcacheFdCacheOverloaded	Long	
rfcacheFdMonitorErrorStuckIo	Long	
rfcachePoolNumSrcDevs	Long	
rfcachePoolNumCacheDevs	Long	
rfcachePoolSize	Long	
rfcachePoolReadHit	Long	
rfcachePoolReadMiss	Long	
rfcachePoolWriteHit	Long	
rfcachePoolWriteMiss	Long	
rfcachePoolCachePages	Long	
rfcachePoolPagesInuse	Long	
rfcachePoolEvictions	Long	

Table 41 System Statistics Object (continued)

Properties	Type	Note
rfcachePoolInLowMemoryCondition	Long	
rfcachePoolIoTimeGreater1Min	Long	
rfcachePoolLockTimeGreater1Sec	Long	
rfcachePoolSuspendedIos	Long	
rfcachePoolLowResourcesInitiatedPassthroughMode	Long	
rfcachePoolosOutstanding	Long	
rfcachePoolReadsPending	Long	
rfcachePoolWritePending	Long	
rfcachePoolSuspendedPequestsRedundantSearchs	Long	
rfcachePoolReadPendingG1Sec	Long	
rfcachePoolReadPendingG10Millis	Long	
rfcachePoolReadPendingG1Millis	Long	
rfcachePoolReadPendingG500Micro	Long	
rfcachePoolWritePendingG1Sec	Long	
rfcachePoolWritePendingG10Millis	Long	
rfcachePoolWritePendingG1Millis	Long	
rfcachePoolWritePendingG500Micro	Long	
rfcachePoolSourceIdMismatch	Long	

Table 42 System Bandwidth Counters

Properties	Type
totalWeightInKb	Long
numOccured	Long
numSeconds	Long

Table 43 System Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/System/instances	GET			List of all System objects
	/api/instances/System::{id}	GET			System object
	/api/instances/System::{id}/relationships/ProtectionDomain	GET			List of Protection Domain objects that are connected to the System
	/api/instances/System::{id}/relationships/Sdc	GET			List of SDC objects that are connected to the System
	/api/instances/System::{id}/relationships/User	GET			List of User objects that are connected to the System
	/api/instances/System::{id}/relationships/Statistics	GET			List of Statistics for the System

	/api/types/System/instances/action/querySelectedStatistics	POST	Required:
			`properties` - list of properties to fetch (to see available properties look at Statistics)
			properties example:
			"properties":["failedCapacityInKb ","movingCapacityliKb "]
					All selected properties
					`{"ids": ["a7e798d100000000"], "properties": ["failedCapacityInKb","movingCapacityIiKb">`
					Response:
					`{ ""a7e798d100000000":{ o "failedCapacityInKb":1, o`

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
					`"movingCapacityIiKb":0`
	/api/types/System/instances/action/queryIdByKey	POST			id

	/api/instances/System::{id}/action/setCapacityAlertThresholds	POST	Required:
			`capacityAlertHighThresholdPercent`
			AND\OR
			`capacityAlertCriticalThresholdPercent`
			`allStoragePools` (TRUE \FALSE)
			AND\OR
			`systemDefault`(TRUE \FALSE)
			At least one of the parameters `allStoragePools` or `systemDefault` must be set to TRUE
				SystemValue sets the default value for the newly created StoragePool	

	/api/instances/System/action/createMdmCluster	POST	Required: `mdmIp` - IP address of the MDM actor that will run the request `ips` - list of MDM IP addresses of the created MDM Optional:	If "mdmAddresses" is defined in the rest gateway configuration, mdmIp should be one of the IP addresses in it. Ensure that each server	id

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			`managementIps` - list of management IP addresses of the created MDM `name` - name of the created MDM `virtIpIntfs` - NIC interface(s) to be used by master MDM for virtual IP addresses `virtIPs` – virtual IP addresses that will be used by cluster `updateConfiguration` - TRUE or FALSE. If set to TRUE and "mdmAddresses" in the REST gateway configuration is empty, "mdmAddresses" will be populated according to the values in the request.	on which MDM is installed has a unique host name. If the error message SERVER_CERTIFICATE_NOT_TRUSTEDis is returned, the MDM's host name is not a unique name. To solve the problem, change the MDM host's name.	
	/api/instances/System/action/addPrimaryMdm	POST	Required: `mdmIp` – IP address of the primary MDM, `primaryMdmActorIpList` – list of IP addresses Optional:	Deprecate (not working from Version 2.0). If “mdmAddresses” is defined in the ScaleIO Gateway configuration	id

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			`mdmManagementIps` – list of IP addresses, `updateConfiguration` – TRUE or FALSE. If set to TRUE and “mdmAddresses” in ScaleIO Gateway configuration is empty, “mdmAddresses” will be populated according to the values in the request.	, mdmIp should be one of the IP addresses in it.	
	/api/instances/System/action/addStandbyMdm	POST	Required: `ips`- list of IP addresses `role` - Mdm type. Manager or TieBreaker Optional: `managementIps`- list of IP addresses `name` - MDM Name `virtIpIntfs` - list of NIC interfaces that will be used for the purpose of virtual IP address `port` - MDM actor port `forceClean` Clean a previous SDS configuration - TRUE or FALSE		id

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			`allowAsymmetricIps`-Allow the added node to have a different number of IP addresses from the Master. TRUE or FALSE		
	/api/instances/System/action/addSecondaryMdm	POST	Optional: `secondaryMdmActorIpList`, `secondaryMdmActorPort`	Deprecate (not working from version 2.0)	
	/api/instances/System/action/addTiebreaker	POST	Optional: `tiebreakerMdmIpList`, `tiebreakerMdmPort`, `forceClean` Clean a previous SDS configuration – TRUE or FALSE. TRUE: Clean previous configuration of the Tie-Breaker	Deprecate (not working from version 2.0)	
	/api/instances/System/action/removeStandbyMdm	POST	Required: `id` - MDM ID (unsigned long, in hexadecimal presentation, for example: 1234abcd)		
	/api/instances/System/action/removeSecondaryMdm	POST		Deprecate (not working from version 2.0)	

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/System/action/removeTiebreaker	POST		Deprecate (not working from version 2.0)	
	/api/instances/System/action/setManagementIps	POST	Required from version 2.0: `id` - MDM ID (unsigned long, in hexadecimal presentation, for example: 1234abcd) Required: `clearAll` - "TRUE" or "FALSE" `mdmManagementIpList` - list of IP addresses Optional: `allowDuplicate` - TRUE or FALSE. If the flag is set, the command will not fail when there is duplication of IP addresses in different MDMs after the modification.	For example: {"id":"76ba26851a1bd2a2","clearAll":"FALSE", "mdmManagementIpList": ["10.76.60.150", "10.76.60.168">	
	/api/instances/System/action/changeMdmOwnership	POST	Required from version 2.0: `id` - new master MDM ID (unsigned long, in hexadecimal presentation, for example: 1234abcd)		
	/api/instances/	POST		Deprecate (not working	

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	System/action/switchMdmOwnership			from version 2.0. use changeMdmOwnership instead)	
	/api/instances/System/action/switchClusterMode	POST	Required: `mode` - new cluster mode (OneNode, ThreeNodes or FiveNodes) `addSlaveMdmIdList` - list of standby slaves to add to cluster `addTBIdList` - list of standby tie breaker IDs to add to cluster Or `removeSlaveMdmIdList` - list of slaves IDs to remove from cluster `removeTBIdList` - list of tie breakers ids to remove from cluster (ID is in hexadecimal presentation, for example: 1234abcd) Optional: `allowLeaveFailed` - "TRUE" or "FALSE" Allow leaving currently failed MDMs in the cluster.	If the cluster is expanded, slave and tie breakers will be added, according to the expanded size. If the cluster size is reduced, slave and tie breakers will be removed.	

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/System/action/replaceClusterMdm	POST	Required: (`addSlaveMdmIdList` - list of standby slaves to add to cluster and `removeSlaveMdmIdList` - list of slave IDs to remove from cluster) Or (`addTBIdList` - list of standby tie breaker IDs to remove from cluster and `removeTBIdList` - list of tie breaker IDs to remove from cluster) (ID is in hexadecimal presentation, for example: 1234abcd) Optional: `allowLeaveFailed` - "TRUE" or "FALSE" Allow leaving currently failed MDMs in the cluster.	Up to half (rounded down) the cluster size can be changed at a time.	
	/api/instances/System/action/startUpgrade	POST			
	/api/instances/System/	POST			

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	action/ abortUpgrade				
	/api/ instances/ System/ action/ finalizeUpgrade	POST			
	/api/ instances/ System/ action/ allowCommandsDuringUpgrade	POST			
	/api/ instances/ System/ action/ enterMaintenanceMode	POST	Required: `sdsId` - SDS ID (unsigned long, in hexadecimal presentation, for example: 1234abcd) or `fsId` - fault set ID (unsigned long, in hexadecimal presentation, for example: 1234abcd) Optional: `forceInsufficientSpace` - TRUE or FALSE `forceDegradedOrFailed` - TRUE or FALSE		List of SDS IDs with command result code. Example: {"sdsMMResults": [{"message":"TGT_ALREADY_IN_MM","resultCode":"425","sdsId":"2046c 61700000002"}>
	/api/ instances/ System/ action/ exitMaintenanceMode i	POST	Required: `sdsId` - SDS ID (unsigned long, in hexadecimal presentation, for example: 1234abcd) or		List of SDS IDs with command result code. Example: {"sdsMMResults": [{"message":"TG

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			`fsId` - fault set ID (unsigned long, in hexadecimal presentation, for example: 1234abcd) Optional: `forceFailedSds` - TRUE or FALSE `forceFailedDevIds` - list of device ids Example: ["a431aa5500010002","a431aa5000010001"]		T_ALREADY_IN_MM","resultCode":"425","sdsId":"2046c 61700000002"}>
	/api/instances/System/action/setCliPasswordAcceptance	POST	Required: `cliPasswordAllowed` - TRUE or FALSE		
	/api/instances/System/action/SetGuidDisplayMode	POST	Required: `showGuid`- TRUE or FALSE		
	/api/instances/System/action/setAuthenticationMethod	POST	Required: `authenticationMethod` - Native, Ldap or NativeAndLdap	“Native” represents local authentication.	
	/api/instances/System/action/generateMdmCsrFile	POST	Required: `Id` - MDM ID (unsigned long, in hexadecimal presentation, for example: 1234abcd) Optional:		

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			`commonName`- Common Name to appear in the CSR. Default is the machine's hostname `country`- Country to appear in the CSR.Max length 2 characters. `state` - State to appear in the CSR `location` - Location to appear in the CSR `organization` - Organization to appear in the CSR `organizationalUnit` - Organizational Unit to appear in the CSR `emailAddress` - Email Address to appear in the CSR		
	/api/instances/System/action/generateMdmCertificate	POST	Required: `Id` - MDM ID (unsigned long, in hexadecimal presentation, for example: 1234abcd)		
	/api/instances/System/action/replaceMdmSecurityFiles	POST	Required: `Id` - MDM ID (unsigned long, in hexadecimal presentation, for example: 1234abcd)		

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/System/action/refreshMdmClusterCapabilities	POST			
	/api/instances/System/action/setConnectionSecurityPolicy	POST	Required: `mdmToSdsPolicy` - None or Authentication		
	/api/instances/System/action/setMdmMode	POST	Required: `mdmMode` – single or cluster	Deprecate (not working from version 2.0)	

	/api/instances/System::{id}/action/setLicense	POST	Required:
			`files - file containing the license key.`
			`Content-type should be multipart/form-data`
				Example using curls
				curl -k -uadmin: $token -X POST -H "Content-Type:multipart/form-data" -F "files=@c:/4b6b488c6656f5a5.key" https://localhost/api/instances/System::f69dbc5b00000000/action/setLicense	

	/api/instances/System::{id}/action/removeConsistencyGroupSnapshots	POST	Required: `snapGroupId`		numberOfVolumes – number of volumes that were removed because of this operation
	/api/instances/	POST	Required:		volumeIdList snapshotGroupId

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return

	System::{id}/action/snapshotVolumes		`snapshotDefs` – a list of combination of “volumeId” volume ID and “snapshotName” (optional field) snapshot name.
			For example: { “snapshotDefs”: [{“ volumeId”: ”2dd9132300000000”, “snapshotName”:”snap1”}, {“volumeId”:”1234”}>
					for example:
					`{`
					`"volumeIdList":[`
					`"2dd9132400000001"`
					`],`
					`"snapshotGroupId":"d2e53daf00000001"`
					`}`

	/api/instances/System::{id}/action/setDrlMode	POST	Required: `drlMode` - "Volatile" or "NonVolatile"		
	/api/instances/System::{id}/action/setObfuscationProperties	POST	Required: `isObfuscated`		
	/api/instances/System::{id}/action/setSystemName	POST	Required: `newName`		

	/api/instances/System::{id}/action/startRemoteSyslog	POST	Required:
			`hostIp - String. IP address of the syslog server.`
			`Can be more than one address, separated by a comma.`
			Optional:		

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			hostPort, facility, attachEventCode - TRUE or FALSE		
	/api/instances/System::{id}/action/setSyslogFacility	POST	Required: facility hostIp-String. IP address of the syslog server. Can be more than one address, separated by a comma.		
	/api/instances/System::{id}/action/stopRemoteSyslog	POST	Required: hostIp—String. IP address of the syslog server. Can be more than one address, separated by a comma.		
	/api/instances/System::{id}/action/setRestrictedSdcModeEnabled	POST	Required: restrictedSdcModeEnabled - TRUE or FALSE		
	/api/instances/System/queryMdmCluster	POST		Doesn't require username and password	
	/api/instances/System/action/	POST			

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	enableManagementClientSecureCommunication				
	/api/instances/System/action/setSdsPerformanceParameters	POST	Required: `perfProfile` - "Default" or "HighPerformance		
	/api/instances/System/action/setSdcPerformanceParameters		Required: `perfProfile` - "Default" or "HighPerformance"		
	/api/instances/System/action/setMdmPerformanceParameters	POST	Required: `perfProfile` - "Default" or "HighPerformance"		
	/api/instances/System/action/resetSdsOscillatingCounter	POST	Required: `allCounters` - TRUE or FALSE OR `failureCounter` - one of the following: `MdmSdsNetworkDisconnections`, `SdsSdsNetworkDisconnections`, `SdsDecoupled`, `SdsConfigurationFailures`, `SdsReceiveBu`		

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			`fferAllocationFailures`		
	/api/instances/System/action/setSdsOscillatingCounterParameters	POST	Required: `failureCounter` – see previous operation for possible values `windowType` – Short or Medium or Long `windowInterval` – number between 1 and 86400 `failuresThreshold`	Set counter parameters to all SDSs	
	/api/instances/System/action/resetSdcOscillatingCounter	POST	Required: `allCounters` – TRUE or FALSE OR `failureCounter` – one of the following: `SdcSdsNetworkDisconnections`, `SdcMdmNetworkDisconnections`, `SdcMemoryAllocationFailures`, `SdcSocketAllocationFailures`, `SdcLongOperations`		
	/api/instances/System/action/setSdcOscillati	POST	Required: `failureCounter` - see previous operation for possible values	Set counter parameters to all SDCs	

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	ngCounterParameters		`windowType` - Short, Medium or Long `windowInterval` - number between 1 and 86400 `failuresThreshold`		
	/api/instances/System/action/resetDeviceOscillatingCounter	POST	Required: `allCounters` - TRUE or FALSE OR `failureCounter` - one of the following: DevLongSuccessfullos		
	/api/instances/System/action/setDeviceOscillatingCounterParameters	POST	Required: `failureCounter` - see previous operation for possible values `windowType` - Short, Medium or Long `windowInterval` - number between 1 and 86400 `failuresThreshold`	Set counter parameters to all Devices	
	/api/instances/System/action/setMdmIps	POST	Required: mdmIpsList - list of ips		
Add LDAP Service	/api/instances/System/action/	POST	Required: `ldapServiceURI`- URI of the LDAP service.		ldapServiceI

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	addLdapService		<LDAP_SCHEMA>:// <LDAP_HOSTNAME>:<PORT_NUMBER > Where <LDAP_SCHEMA> defines the connection protocol: LDAPS - secure LDAP connection (recommended), LDAP - non-secured TCP LDAP connection <PORT_NUMBER> - An optional parameter of LDAP service port `ldapBaseDn`- Base DN (Distinguished Name) of users in domain Example: {"ldapServiceURI":"ldaps://domain.name.ldaps.local", "ldapBaseDn":"OU=SIO_OU_1,DC=ldaps,DC=local"} Optional: `ldapServiceName` -Name of LDAP service `ldapObjectClassAttr`- The object class attribute that is		

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			used to identify a user. Used in the search filter (default is: 'user')		
			`ldapUserIdAttrName`-The attribute name that defines the user ID used in the search filter (default is: 'sAMAccountName') `ldapMemberOfAttrName`- The attribute name that defines the contained group used in the search filter (default is: 'memberOf') `disable_recursive_search` -TRUE or FALSE. Disables recursive search		
Assign group to ldap service	/api/instances/System/action/assignRoleToLdapGroup	POST	Required: `ldapServiceId` roleToGroupAssignment (List of RoleToLdapGroupAssignment) `bOverwrite` [true,false] Example: {"ldapServiceId": "cfc9463400000000","roleToGroupAssignment": [{"role": "Administrator", "groupDn":	Role can be one of the following parameters: Monitor, Configure, Administrator, Security, FrontendConfig or BackendConfig	

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			"CN=SIO_GRP_1,OU=SIO_OU_1,DC=ldaps,DC=local"}],"bOverwrite": "true"}		
Remove LDAP service	/api/instances/System/action/deleteLdapService	POST	Required: `ldapServiceId` OR `allLdapServices` —with empty value Example: {"ldapServiceId": "cfc9463400000000",""} or {"allLdapServices": ""}		
Query LDAP Service	/api/instances/System/action/queryLdap	POST			List of Maps of service properties: ldapServiceURI, ldapBaseDn, ldapserviceId, ldapSearchFilter, ldapGroupsMappingList Example: {"ldapServicesList": [{"ldapServiceURI":"ldaps://win12r2-dc.ldaps.local","ldapGro upsMappingList": "[{role=Administrator, groupDn=CN=SIO_GRP_1,OU=SIO_OU_1,DC=ldaps,DC=local}]","ldapServiceId":"bebb848ff00000000

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
					","ldapSearchFilter":"(&(objectClass=posixAccount)(uid=Administrator)(memberOf=SIO_GRP_1))","ldapBaseDn":"ou=sio_ou_1,dc=ldaps,dc=local"}>

Rename LDAP Service	/api/instances/System/action/renameLdapService	POST	Required:
			`ldapServiceId`
			`newName`
			Example:
			{"ldapServiceId": "cfc9463400000000","newName": "TestService"}		

Remove LDAP Group	/api/instances/System/action/removeLdapGroup	POST	Required: `ldapServiceId` `roleList` - list of user roles. Possible roles: Monitor, Configure, Administrator, Security, FrontendConfig, BackendConfig		
	/api/instances/System/action/querySdsOscillatingNetworkCounters	POST			
	/api/instances/System/action/querySdcOscil	POST			

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	latingNetwork Counters				
	/api/instances/System/action/disableAdmin	POST			
	/api/instances/System/action/renameMdm	POST	Required: `Id -mdm id` in hexadecimal presentation. For example: 1234abcd newName		
Modify MDM virtual IP interfaces	/api/instances/System/action/modifyVirtualIpInterfaces	POST	Required: `Id` - mdm manager id in hexadecimal presentation. For example: 5ce68c755220ad30 And one of the following: `virtIpIntfs` - list of NIC interfaces to be used for virtual IPs `clear` - Boolean indicating whether to clear existing NIC interfaces list		

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
Modify system virtual (floating) IP addresses	/api/instances/System/action/modifyClusterVirtualIPs	POST	One of the following: `virtIPs` - list of virtual IPs to be configured for the system `clear` - Boolean indicating whether to clear existing NIC interfaces list	**Note** SDC settings should be modified manually after using this URI. Failure to do so may cause IO errors, due to SDC communication errors.	
Query system limits	/api/instances/System/action/querySystemLimits	POST			List of Maps of global system limits. Map fields: type, description, minVal, maxVal Possible types: protectionDomains,storagePools,faultSets,sdss,sdcs,vtrees,volumes,snapshots,users,ldapEntries,sessions,rsyslogClients,storagePoolsPerProtectionDomain,sdssPerProtectionDomain,faultSetsPerProtectionDomain,sdssPerFaultSet,devicesPerSds,rfcacheDevicePerSds,devicePerStoragePool,volumePerVtree,volumesPerProtectionDomain,snapshotsPerVtree,volumesPerSdc,ldapGroupsPerLdapEntry,ipPerMdm,ipPerHost,volumeSizeGb,allVolumesSizeGb

Table 43 System Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
					Example of a system limit info: {"maxVal":"1048576","minVal":"8","type":"volumeSizeGb","description":"Volume size in GB"}

Protection Domain

Table 44 Protection Domain Object

Properties	Type	Required\Optional for POST
id	Name	
name	Name	Required
protectionDomainState	Active, ActivePending, Inactive or InactivePending	-
rebuildNetworkThrottlingEnabled	Boolean	-
rebalanceNetworkThrottlingEnabled	Boolean	-
overallIoNetworkThrottlingEnabled	Boolean	-
rebuildNetworkThrottlingInKbps	Integer	-
rebalanceNetworkThrottlingInKbps	Integer	-
overallIoNetworkThrottlingInKbps	Integer	-
systemId	String	-
sdsDecoupledCounterParameters	OscillatingCounterParameters	
sdsConfigurationFailureCounterParameters	OscillatingCounterParameters	
mdmSdsNetworkDisconnectionsCounterParameters	OscillatingCounterParameters	
sdsSdsNetworkDisconnectionsCounterParameters	OscillatingCounterParameters	
sdsReceiveBufferAllocationFailuresCounterParameters	OscillatingCounterParameters	

Table 44 Protection Domain Object (continued)

Properties	Type	Required\Optional for POST
rfcacheEnabled	Boolean	
rfcachePageSizeKb	Integer	
rfcacheMaxIoSizeKb	Integer	
rfcacheOpertionalMode	Read, Write, ReadAndWrite, WriteMiss	

Table 45 Protection Domain Parent

Name	Type	Note
System	System	

Table 46 Protection Domain Relationships

Name	Type	Note
SDS	List of SDS objects	
StoragePool	List of StoragePool objects	
FaultSet	List of Fault Set objects	
Statistics		

Protection Domain Statistics and Operations

Table 47 Protection Domain Statistics Object

Properties	Type	Note
numOfSds	Long	
numOfStoragePools	Long	
numOfFaultSets	Long	
rmcacheSizeInKb	Long	
capacityLimitInKb	Long	
maxCapacityInKb	Long	
capacityInUseInKb	Long	
thickCapacityInUseInKb	Long	
thinCapacityInUseInKb	Long	

Table 47 Protection Domain Statistics Object (continued)

Properties	Type	Note
snapCapacityInUseInKb	Long	
snapCapacityInUseOccupiedInKb	Long	
unreachableUnusedCapacityInKb	Long	
protectedVacInKb	Long	
degradedHealthyVacInKb	Long	
degradedFailedVacInKb	Long	
failedVacInKb	Long	
inUseVacInKb	Long	
activeMovingInFwdRebuildJobs	Long	
pendingMovingInFwdRebuildJobs	Long	
activeMovingOutFwdRebuildJobs	Long	
pendingMovingOutFwdRebuildJobs	Long	
activeMovingInBckRebuildJobs	Long	
pendingMovingInBckRebuildJobs	Long	
activeMovingOutBckRebuildJobs	Long	
pendingMovingOutBckRebuildJobs	Long	
activeMovingInRebalanceJobs	Long	
pendingMovingInRebalanceJobs	Long	
activeMovingRebalanceJobs	Long	
pendingMovingRebalanceJobs	Long	
primaryVacInKb	Long	“primary” represents the Master MDM in version 2.0 and higher
secondaryVacInKb	Long	“secondary” represents the Slave MDM in version 2.0 and higher
primaryReadBwc	Bandwidth Counter	“primary” represents the Master MDM in version 2.0 and higher
primaryReadFromDevBwc	Bandwidth Counter	“primary” represents the Master MDM in version 2.0 and higher
primaryWriteBwc	Bandwidth Counter	“primary” represents the Master MDM in version 2.0 and higher
secondaryReadBwc	Bandwidth Counter	“secondary” represents the Slave MDM in version 2.0 and higher

Table 47 Protection Domain Statistics Object (continued)

Properties	Type	Note
secondaryReadFromDevBwc	Bandwidth Counter	“secondary” represents the Slave MDM in version 2.0 and higher
secondaryWriteBwc	Bandwidth Counter	“secondary” represents the Slave MDM in version 2.0 and higher
totalReadBwc	Bandwidth Counter	
totalWriteBwc	Bandwidth Counter	
fwdRebuildReadBwc	Bandwidth Counter	
fwdRebuildWriteBwc	Bandwidth Counter	
bckRebuildReadBwc	Bandwidth Counter	
bckRebuildWriteBwc	Bandwidth Counter	
rebalanceReadBwc	Bandwidth Counter	
rebalanceWriteBwc	Bandwidth Counter	
spareCapacityInKb	Long	
capacityAvailableForVolumeAllocationInKb	Long	
protectedCapacityInKb	Long	
degradedHealthyCapacityInKb	Long	
degradedFailedCapacityInKb	Long	
failedCapacityInKb	Long	
movingCapacityInKb	Long	
activeMovingCapacityInKb	Long	
pendingMovingCapacityInKb	Long	
fwdRebuildCapacityInKb	Long	
activeFwdRebuildCapacityInKb	Long	
pendingFwdRebuildCapacityInKb	Long	
bckRebuildCapacityInKb	Long	

Table 47 Protection Domain Statistics Object (continued)

Properties	Type	Note
activeBckRebuildCapacityInKb	Long	
pendingBckRebuildCapacityInKb	Long	
rebalanceCapacityInKb	Long	
activeRebalanceCapacityInKb	Long	
pendingRebalanceCapacityInKb	Long	
atRestCapacityInKb	Long	
numOfUnmappedVolumes	Long	
numOfMappedToAllVolumes	Long	
numOfThickBaseVolumes	Long	
numOfThinBaseVolumes	Long	
numOfSnapshots	Long	
numOfVolumesInDeletion	Long	
fixedReadErrorCount	Long	
BackgroundScanCompareCount	Long	
BackgroundScannedInMB	Long	
thinCapacityAllocatedInKm	Long	
rmPendingAllocatedInKb	Long	
semiProtectedVacInKb	Long	
inMaintenanceVacInKb	Long	
activeMovingInNormRebuildJobs	Long	
activeMovingOutNormRebuildJobs	Long	
pendingMovingInNormRebuildJobs	Long	
pendingMovingOutNormrebuildJobs	Long	
primaryReadFromRmcacheBwc	Bandwidth Counter	“primary” represents the Master MDM in version 2.0 and higher
secondaryReadFromRmcacheBwc	Bandwidth Counter	“secondary” represents the Slave MDM in version 2.0 and higher
normRebuildReadBwc	Bandwidth Counter	
normRebuildWriteBwc	Bandwidth Counter	

Table 47 Protection Domain Statistics Object (continued)

Properties	Type	Note
rfcacheReadsReceived	Long	
rfcacheWritesReceived	Long	
rfcacheAvgReadTime	Long	
rfcacheAvgWriteTime	Long	
rfcacheSourceDeviceReads	Long	
rfcacheSourceDeviceWrites	Long	
rfacheReadHit	Long	
rfcacheReadMiss	Long	
rfcacheWriteMiss	Long	
rfcacheIosSkipped	Long	
rfcacheReadsSkipped	Long	
rfcacheReadsSkippedAlignedSizeTooLarge	Long	
rfcacheReadsSkippedMaxIoSize	Long	
rfcacheReadsSkippedHeavyLoad	Long	
rfcacheReadsSkippedStuckIo	Long	
rfcacheReadsSkippedLowResources	Long	
rfcacheReadsSkippedInternalError	Long	
rfcacheReadsSkippedLockIos	Long	
rfcacheWritesSkippedMaxIoSize	Long	
rfcacheWritesSkippedHeavyLoad	Long	
rfcacheWritesSkippedStuckIo	Long	
rfcacheWritesSkippedLowResources	Long	
rfcacheWritesSkippedInternalError	Long	
rfcacheWritesSkippedCacheMiss	Long	
rfcacheSkippedUnlinedWrite	Long	
rfcacheIoErrors	Long	
rfcacheReadsFromCache	Long	
rfcacheIosOutstanding	Long	
rfcacheReadsPending	Long	

Table 47 Protection Domain Statistics Object (continued)

Properties	Type	Note
rfcacheWritePending	Long	
semiProtectedCapacityInKb	Long	
inMaintenanceCapacityInKb	Long	
normRebuildCapacityInKb	Long	
activeNormRebuildCapacityInKb	Long	
pendingNormRebuildCapacityInKb	Long	
rfcacheFdReadsReceived	Long	
rfcacheFdWritesReceived	Long	
rfcacheFdInlightReads	Long	
rfcacheFdInlightWrites	Long	
rfcacheFdReadTimeGreater500Millis	Long	
rfcacheFdReadTimeGreater1Sec	Long	
rfcacheFdReadTimeGreater5Sec	Long	
rfcacheFdReadTimeGreater1Min	Long	
rfcacheFdWriteTimeGreater500Millis	Long	
rfcacheFdWriteTimeGreater1Sec	Long	
rfcacheFdWriteTimeGreater5Sec	Long	
rfcacheFdWriteTimeGreater1Min	Long	
rfcacheFdAvgReadTime	Long	
rfcacheFdAvgWriteTime	Long	
rfcacheFdIoErrors	Long	
rfcacheFdCacheOverloaded	Long	
rfcacheFdMonitorErrorStuckIo	Long	
rfcachePoolNumSrcDevs	Long	
rfcachePoolNumCacheDevs	Long	
rfcachePoolSize	Long	
rfcachePoolReadHit	Long	
rfcachePoolReadMiss	Long	
rfcachePoolWriteHit	Long	

Table 47 Protection Domain Statistics Object (continued)

Properties	Type	Note
rfcachePoolWriteMiss	Long	
rfcachePoolCachePages	Long	
rfcachePoolPagesInuse	Long	
rfcachePoolEvictions	Long	
rfcachePoolInLowMemoryCondition	Long	
rfcachePoolIoTimeGreater1Min	Long	
rfcachePoolLockTimeGreater1Sec	Long	
rfcachePoolSuspendedIos	Long	
rfcachePoolLowResourcesInitiatedPassthroughMode	Long	
rfcachePoolosOutstanding	Long	
rfcachePoolReadsPending	Long	
rfcachePoolWritePending	Long	
rfcachePoolSuspendedPequestsRedundantSearchs	Long	
rfcachePoolReadPendingG1Sec	Long	
rfcachePoolReadPendingG10Millis	Long	
rfcachePoolReadPendingG1Millis	Long	
rfcachePoolReadPendingG500Micro	Long	
rfcachePoolWritePendingG1Sec	Long	
rfcachePoolWritePendingG10Millis	Long	
rfcachePoolWritePendingG1Millis	Long	
rfcachePoolWritePendingG500Micro	Long	
rfcachePoolSourceIdMismatch	Long	

Table 48 Protection Domain Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/ProtectionDomain/instances	GET			List of all Protection Domain objects

Table 48 Protection Domain Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/ProtectionDomain::{id}	GET			Protection Domain object
	/api/instances/ProtectionDomain::{id}/relationships/Sds	GET			List of SDS objects that are connected to the Protection Domain
	/api/instances/ProtectionDomain::{id}/relationships/StoragePool	GET			List of StoragePool objects that are connected to the Protection Domain
	/api/instances/ProtectionDomain::{id}/relationships/FaultSet	GET			List of Fault Set objects that are connected to the Protection Domain
	/api/instances/ProtectionDomain::{id}/relationships/Statistics	GET			List of Statistics for the Protection Domain
	/api/types/ProtectionDomain/instances/action/querySelectedStatistics	POST	Required: `ids` - list of objects IDs or `allIds` - with empty value properties - list of properties to fetch (to see available properties look at Statistics)	If `allIds` appears, ids list is ignored and a list of selected properties is returned for every ProtectionDomain.	List of: id and the selected properties. For example: https://localhost:443/api/types/ProtectionDomain/instances/action/querySelectedStatistics with: {"properties": ["capacityInUseInKb"], "allIds":""} Returns:

Table 48 Protection Domain Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
					{ ""ce475c1800000001":{ o "capacityInUseInKb":0 }, ""ce475c1700000000":{ o "capacityInUseInKb":33554432 } }
	/api/types/ProtectionDomain/instances/action/queryIdByKey	POST	Required: `name`		id
	/api/types/ProtectionDomain/instances/action/queryBySelectedIds	POST	Required: `ids`		List of ProtectionDomain objects matching the ids See the example immediately after this table.
Remove a Protection Domain	/api/instances/ProtectionDomain::{id}/action/removeProtectionDomain	POST			
	/api/types/ProtectionDomain/instances	POST	According to “Required \Optional for POST” column in Protection Domain object.		id
	/api/instances/ProtectionDomain::{id}/	POST	Required: `name`		

Table 48 Protection Domain Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	action/setProtectionDomainName				
	/api/instances/ProtectionDomain::{id}/action/activateProtectionDomain	POST	Required: `forceActivate`		
	/api/instances/ProtectionDomain::{id}/action/inactivateProtectionDomain	POST	Required: `forceShutdown` - "TRUE" or "FALSE"		
	/api/instances/ProtectionDomain::{id}/action/queryDevicesTestResults	POST			Response: List of DeviceTestResults
Limit the network bandwidth of all SDSs in the Protection Domain for various traffic types	/api/instances/ProtectionDomain::{id}/action/setSdsNetworkLimits	POST	Required: `rebuildLimitInKbps` AND \OR `rebalanceLimitInKbp` AND\OR `overallLimitInKbps`	Overall limit size must be larger than both rebuild and rebalance, and all limits must be more than 5 MB	
	/api/instances/ProtectionDomain::{id}/action/queryProtectionDomainNetwork	POST			Response: List of SdsPeerConnectionState
	/api/instances/ProtectionDomain::{id}/	POST	Required: `rmcacheEnabled` - TRUE or FALSE		

Table 48 Protection Domain Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	action/setSdsRmcacheEnabled				
	/api/instances/ProtectionDomain::{id}/action/setSdsRmcacheSize	POST	Required: `rmcacheSizeInMB`		
	/api/instances/ProtectionDomain::{id} /action/setRfcacheParameters	POST	Required: `pageSizeKb`—Cache page size in KB. Valid range: 4-64 AND\OR `maxIOSizeKb`—Cache maximum IO size in KB. Valid range: 32-256 AND\OR `rfcacheOperationMode`—Pass-through mode. One of: "None" or "Read" or "Write" or "ReadAndWrite" or "WriteMiss"		
	/api/instances/ProtectionDomain::{id}/action/enableSdsRfcache	POST			
	/api/instances/ProtectionDomain::{id}/action/	POST			

Table 48 Protection Domain Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	disableSdsRfcache				
	/api/instances/ProtectionDomain::{id}/action/setSdsPerformanceParameters	POST	Required: `perfProfile` - "Default" or "HighPerformance"		
	/api/instances/ProtectionDomain::{id}/action/resetSdsOscillatingCounter	POST	Required: `allCounters` - TRUE or FALSE OR `failureCounter` - one of the following: MdmSdsNetworkDisconnections, SdsSdsNetworkDisconnections, SdsDecoupled, SdsConfigurationFailures,SdsReceiveBufferAllocationFailures		
	/api/instances/ProtectionDomain::{id}/action/setSdsOscillatingCounterParameters	POST	Required: `failureCounter` - see previous operation for possible values `windowType` - Short or Medium or Long `windowInterval` - number between 1 and 86400 failuresThreshold		

Example for `queryBySelectedIds` (ProtectionDomain objects matching the IDs):

```
 /api/types/ProtectionDomain/instances/action/queryBySelectedIds
Body: {"ids":["2c3071a200000001">
Response:
[
{
"name":"Lab2",
"systemId":"57b89a5b22b0bc80",
"protectionDomainState":"Active",
"rebuildNetworkThrottlingInKbps":null,
"rebalanceNetworkThrottlingInKbps":null,
"overallIoNetworkThrottlingInKbps":null,
"rebuildNetworkThrottlingEnabled":false,
"rebalanceNetworkThrottlingEnabled":false,
"overallIoNetworkThrottlingEnabled":false,
"id":"2c3071a200000001",
"links":[
{
"rel":"self",
"href":"/api/instances/ProtectionDomain::2c3071a200000001"
},
{
"rel":"/api/ProtectionDomain/relationship/Statistics",
"href":"/api/instances/ProtectionDomain::2c3071a200000001/
relationships/Statistics"
},
{
"rel":"/api/ProtectionDomain/relationship/FaultSet",
"href":"/api/instances/ProtectionDomain::2c3071a200000001/
relationships/FaultSet"
},
{
"rel":"/api/ProtectionDomain/relationship/StoragePool",
"href":"/api/instances/ProtectionDomain::2c3071a200000001/
relationships/StoragePool"
},
{
"rel":"/api/ProtectionDomain/relationship/Sds",
"href":"/api/instances/ProtectionDomain::2c3071a200000001/
relationships/Sds"
},
{
"rel":"/api/parent/relationship/systemId",
"href":"/api/instances/System::57b89a5b22b0bc80"
}
]
}
]
```

Table 49 SdsConnectionState

Properties	Type
sdsId	String
peerConnectionStates	List of SdsPeerConnectionState s

Table 50 SdsPeerConnectionState

Properties	Type
peerSdsId	String

Table 50 SdsPeerConnectionState (continued)

Properties	Type
connections	List of SingleConnectionState s

Table 51 SingleConnectionState

Properties	Type
peerHostname	String
isConnected	Boolean
connectionCount	Integer

Table 52 DeviceTestResults

Properties	Type
deviceId	String
deviceTestResults	List of DeviceTestResult

Table 53 DeviceTestResult

Properties	
testTime	Long
totalIOInKb	Long
numOfIos	Integer
millis	Integer
testNum	Integer
pattern	PATTERN_NONE, PATTERN_WRITE_SEQ, PATTERN_WRITE_RAND, PATTERN_READ_SEQ, PATTERN_READ_RAND or PATTERN_LAST
rc	Integer

SDS

Table 54 SDS Object

Properties	Type	Required\Optional for POST
id	String	-

Table 54 SDS Object (continued)

Properties	Type	Required\Optional for POST
name	String	Optional
sdsIpList	SdsIp list	Required—name of the property for POST: sdsIpList
sdsPort	Integer	Optional—name of the property for POST: sdsPort
sdsState	Normal or RemovePending	-
membershipState	JoinPending or Joined or Decoupled	-
mdmConnectionState	Connected or Disconnected	-
drlMode	Volatile or NonVolatile Hardened(2), MemOnly	Optional
rmcacheEnabled	Boolean	Optional
rmcacheSizeInKb	Integer	Optional
rmcacheFrozen	Boolean	Optional
protectionDomainId	String	Required
onVmWare	Boolean	-
faultSetId	String	Optional
numOfIoBuffers	Integer	Optional
RmcacheMemoryAllocationState	RmcacheMemoryAllocationStateInvalid or AllocationPending or AllocationSuccessful or AllocationFailed or RmcacheDisabled	-
perfParams	SdsPerfParams	
currProfilePerfParam	SdsPerfParams	
maintenanceState	NoMaintenance, SetMaintenanceInProgress, InMaintenance or ExitMaintenanceInProgress	
softwareVersionInfo	String	
rfcacheErrorDeviceDoesNotExis	Boolean	
rfcacheErrorLowResources	Boolean	

Table 54 SDS Object (continued)

Properties	Type	Required\Optional for POST
rfcacheErrorApiVersionMismatch	Boolean	
rfcacheErrorInconsistentCacheConfiguration	Boolean	
rfcacheErrorInconsistentSourceConfiguration	Boolean	
rfcacheErrorInvalidDriverPath	Boolean	
certificateInfo	CertificateInfo	
authenticationError	None, General, ErrorLoadingOpenssl, ErrorLoadingCertificate, VerificationError, ErrorLoadingAuthenticationInMdm, OpensslVerionTooLong	authenticationError is optional
forceDeviceTakeover	Boolean	forceDeviceTakeover is optional
sdsDecoupled	OscillatingCounter	
sdsConfigurationFailure	OscillatingCounter	
sdsReceiveBufferAllocationFailures	OscillatingCounter	

Table 55 SDS Parent

Name	Type	Note
ProtectionDomain	Protection Domain	
FaultSet	Fault Set	

Table 56 SDS Relationships

Name	Type	Note
Device	List of Device objects	
RfcacheDevice	List of RfcacheDevice objects	
Statistics		

Table 57 SdsIp

Properties	Type
ip	String
role	sdsOnly or sdcOnly or all

Table 58 SdsRfcacheDeviceInfo

Properties	Type
path	String
name	String

Table 59 CertificateInfo

Properties	Type
subject	String
issuer	String
validFrom	String
validTo	String
thumbprint	String
validFromAsn1Format	String
validToAsn1Format	String

Table 60 SdsPerfParams

Properties	Type
mdmNumberSocketsPerSdsIp	Integer
mdmSdsKeepaliveTime	Integer
sdsNumberNetworkUmt	Integer
sdsTcpSendBufferSize	Integer
sdsTcpReceiveBufferSize	Integer
sdsMaxNumberAsynchronousIoPerDevice	Integer
sdsNumberSdcIoUmt	Integer
sdsNumberSdsIoUmt	Integer
sdsNumberSdsCopyIoUmt	Integer
sdsNumberCopyUmt	Integer
sdsNumberOsThreads	Integer

Table 60 SdsPerfParams (continued)

Properties	Type
sdsUmtNum	Integer
sdsNumberSocketsPerSdsIp	Integer
sdsNumberIoBuffers	Integer
perfProfile	Default or HighPerformance

SDS Statistics and Operations

Table 61 SDS Statistics Object

Properties	Type	Note
numOfDevices	Long	
rebuildWaitSendQLength	Long	
rebalanceWaitSendQLength	Long	
rebuildPerReceiveJobNetThrottlingInKbps	Long	
rebalancePerReceiveJobNetThrottlingInKbps	Long	
rmcacheSizeInUseInKb	Long	
rmcacheEntryEvictionSizeCountInKb	Long	
rmcacheBigBlockEvictionSizeCountInKb	Long	
rmcacheEntryEvictionCount	Long	
rmcacheBigBlockEvictionCount	Long	
rmcacheNoEvictionCount	Long	
rmcache4kbEntryCount	Long	
rmcache8kbEntryCount	Long	
rmcache16kbEntryCount	Long	
rmcache32kbEntryCount	Long	
rmcache64kbEntryCount	Long	
rmcache128kbEntryCount	Long	
rmcacheCurrNumOf4kbEntries	Long	
rmcacheCurrNumOf8kbEntries	Long	

Table 61 SDS Statistics Object (continued)

Properties	Type	Note
rmcacheCurrNumOf16kbEntries	Long	
rmcacheCurrNumOf32kbEntries	Long	
rmcacheCurrNumOf64kbEntries	Long	
rmcacheCurrNumOf128kbEntries	Long	
rmcacheSkipCountLargeIo	Long	
rmcacheSkipCountUnaligned4kbIo	Long	
rmcacheSkipCountCacheAllBusy	Long	
capacityLimitInKb	Long	
maxCapacityInKb	Long	
capacityInUseInKb	Long	
thickCapacityInUseInKb	Long	
thinCapacityInUseInKb	Long	
snapCapacityInUseInKb	Long	
snapCapacityInUseOccupiedInKb	Long	
unreachableUnusedCapacityInKb	Long	
protectedVacInKb	Long	
degradedHealthyVacInKb	Long	
degradedFailedVacInKb	Long	
failedVacInKb	Long	
inUseVacInKb	Long	
activeMovingInFwdRebuildJobs	Long	
pendingMovingInFwdRebuildJobs	Long	
activeMovingOutFwdRebuildJobs	Long	
pendingMovingOutFwdRebuildJobs	Long	
activeMovingInBckRebuildJobs	Long	
pendingMovingInBckRebuildJobs	Long	
activeMovingOutBckRebuildJobs	Long	

Table 61 SDS Statistics Object (continued)

Properties	Type	Note
pendingMovingOutBckRebuildJobs	Long	
activeMovingInRebalanceJobs	Long	
activeMovingRebalanceJobs	Long	
pendingMovingRebalanceJobs	Long	
primaryVacInKb	Long	“primary” represents Master MDM in versions 2.0 and higher
secondaryVacInKb	Long	“secondary” represents Slave MDM in versions 2.0 and higher
primaryReadBwc	Bandwidth Counter	“primary” represents Master MDM in versions 2.0 and higher
primaryReadFromDevBwc	Bandwidth Counter	“primary” represents Master MDM in versions 2.0 and higher
primaryWriteBwc	Bandwidth Counter	“primary” represents Master MDM in versions 2.0 and higher
secondaryReadBwc	Bandwidth Counter	“secondary” represents Slave MDM in versions 2.0 and higher
secondaryReadFromDevBwc	Bandwidth Counter	“secondary” represents Slave MDM in versions 2.0 and higher
secondaryWriteBwc	Bandwidth Counter	“secondary” represents Slave MDM in versions 2.0 and higher
totalReadBwc	Bandwidth Counter	
totalWriteBwc	Bandwidth Counter	
fwdRebuildReadBwc	Bandwidth Counter	
fwdRebuildWriteBwc	Bandwidth Counter	
bckRebuildReadBwc	Bandwidth Counter	
bckRebuildWriteBwc	Bandwidth Counter	
rebalanceReadBwc	Bandwidth Counter	
rebalanceWriteBwc	Bandwidth Counter	
fixedReadErrorCount	Long	
BackgroundScanCompareCount	Long	
BackgroundScannedInMB	Long	
thinCapacityAllocatedInKm	Long	

Table 61 SDS Statistics Object (continued)

Properties	Type	Note
rmPendingAllocatedInKb	Long	
semiProtectedVacInKb	Long	
inMaintenanceVacInKb	Long	
activeMovingInNormRebuildJobs	Long	
activeMovingOutNormRebuildJobs	Long	
pendingMovingInNormRebuildJobs	Long	
pendingMovingOutNormrebuildJobs	Long	
primaryReadFromRmcacheBwc	Bandwidth Counter	“primary” represents Master MDM in versions 2.0 and higher
secondaryReadFromRmcacheBwc	Bandwidth Counter	“secondary” represents Slave MDM in versions 2.0 and higher
normRebuildReadBwc	Bandwidth Counter	
normRebuildWriteBwc	Bandwidth Counter	
rfcacheReadsReceived	Long	
rfcacheWritesReceived	Long	
rfcacheAvgReadTime	Long	
rfcacheAvgWriteTime	Long	
rfcacheSourceDeviceReads	Long	
rfcacheSourceDeviceWrites	Long	
rfacheReadHit	Long	
rfcacheReadMiss	Long	
rfcacheWriteMiss	Long	
rfcacheIosSkipped	Long	
rfcacheReadsSkipped	Long	
rfcacheReadsSkippedAlignedSizeTooLarge	Long	
rfcacheReadsSkippedMaxIoSize	Long	
rfcacheReadsSkippedHeavyLoad	Long	
rfcacheReadsSkippedStuckIo	Long	

Table 61 SDS Statistics Object (continued)

Properties	Type	Note
rfcacheReadsSkippedLowResources	Long	
rfcacheReadsSkippedInternalError	Long	
rfcacheReadsSkippedLockIos	Long	
rfcacheWritesSkippedMaxIoSize	Long	
rfcacheWritesSkippedHeavyLoad	Long	
rfcacheWritesSkippedStuckIo	Long	
rfcacheWritesSkippedLowResources	Long	
rfcacheWritesSkippedInternalError	Long	
rfcacheWritesSkippedCacheMiss	Long	
rfcacheSkippedUnlinedWrite	Long	
rfcacheIoErrors	Long	
rfcacheReadsFromCache	Long	
rfcacheIosOutstanding	Long	
rfcacheReadsPending	Long	
rfcacheWritePending	Long	
semiProtectedCapacityInKb	Long	
inMaintenanceCapacityInKb	Long	
normRebuildCapacityInKb	Long	
activeNormRebuildCapacityInKb	Long	
pendingNormRebuildCapacityInKb	Long	
rfcacheFdReadsReceived	Long	
rfcacheFdWritesReceived	Long	
rfcacheFdInlightReads	Long	
rfcacheFdInlightWrites	Long	
rfcacheFdReadTimeGreater500Millis	Long	
rfcacheFdReadTimeGreater1Sec	Long	

Table 61 SDS Statistics Object (continued)

Properties	Type	Note
rfcacheFdReadTimeGreater5Sec	Long	
rfcacheFdReadTimeGreater1Min	Long	
rfcacheFdWriteTimeGreater500Millis	Long	
rfcacheFdWriteTimeGreater1Sec	Long	
rfcacheFdWriteTimeGreater5Sec	Long	
rfcacheFdWriteTimeGreater1Min	Long	
rfcacheFdAvgReadTime	Long	
rfcacheFdAvgWriteTime	Long	
rfcacheFdIoErrors	Long	
rfcacheFdCacheOverloaded	Long	
rfcacheFdMonitorErrorStuckIo	Long	
rfcachePoolNumSrcDevs	Long	
rfcachePoolNumCacheDevs	Long	
rfcachePoolSize	Long	
rfcachePoolReadHit	Long	
rfcachePoolReadMiss	Long	
rfcachePoolWriteHit	Long	
rfcachePoolWriteMiss	Long	
rfcachePoolCachePages	Long	
rfcachePoolPagesInuse	Long	
rfcachePoolEvictions	Long	
rfcachePoolInLowMemoryCondition	Long	
rfcachePoolIoTimeGreater1Min	Long	
rfcachePoolLockTimeGreater1Sec	Long	
rfcachePoolSuspendedIos	Long	
rfcachePoolLowResourcesInitiatedPassthroughMode	Long	
rfcachePoolosOutstanding	Long	

Table 61 SDS Statistics Object (continued)

Properties	Type	Note
rfcachePoolReadsPending	Long	
rfcachePoolWritePending	Long	
rfcachePoolSuspendedPequests RedundantSearchs	Long	
rfcachePoolReadPendingG1Sec	Long	
rfcachePoolReadPendingG10Millis	Long	
rfcachePoolReadPendingG1Millis	Long	
rfcachePoolReadPendingG500Micro	Long	
rfcachePoolWritePendingG1Sec	Long	
rfcachePoolWritePendingG10Millis	Long	
rfcachePoolWritePendingG1Millis	Long	
rfcachePoolWritePendingG500 Micro	Long	
rfcachePoolSourceIdMismatch	Long	

Table 62 SDS Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/Sds/instances	GET			List of all SDS objects
	/api/instances/Sds::{id}	GET			SDS object
	/api/instances/Sds::{id}/relationships/Device	GET			List of Devices that are connected to the SDS
	/api/instances/Sds::{id}/relationships/Statistics	GET			List of Statistics for the SDS

Table 62 SDS Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/types/Sds/instances/action/querySelectedStatistics	POST	Required: `ids` - list of object IDs or `allIds` - with empty value `properties` - list of properties to fetch (to see available properties look at Statistics)	If `allIds` appears, ids list is ignored and a list of selected properties is returned for every SDS.	List of id and the selected properties. See the example immediately following this table.
	/api/types/Sds/instances/action/queryIdByKey	POST	Required: `name`		id
	/api/types/Sds/instances/action/queryBySelectedIds	POST	Required: `ids`		List of SDS objects matching the ids
	/api/instances/Sds::{id}/action/removeSds	POST	Optional: force - TRUE or FALSE		
	/api/instances/Sds::{id}/action/abortRemoveSds	POST		This request attempts to abort the removal of a server. If done at a late stage, this request might fail and the server will be removed.	
	/api/instances/Sds::{id}/action/	POST			

Table 62 SDS Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	generateCertificate				
	/api/instances/Sds::{id}/action/queryOscillatingNetworkCounters	POST			
	/api/types/Sds/instances	POST	According to “Required \Optional for POST” column in SDS object. Plus: Optional: `deviceInfoList` – list of: "devicePath", "storagePoolId" and "deviceName" (optional) `sdsRfcacheDeviceInfoList` - list of: "path"and "name" `forceClean` Clean a previous SDS configuration – TRUE or FALSE `deviceTestTimeSecs` - The maximal test run-time in seconds or 128 MB data. Default limit is 10 seconds. Not relevant for noTest.	Parameters syntax example: {"protectionDomainId":"bded100000000","sdsIpList":[{"SdsIp":{"ip":"","role":"all"}}],"deviceInfoList":[{"deviceName":"dev1","devicePath":"/dev/sdb","storagePoolId":"ad65063f00000000"}],"sdsPort":"","name":"SDS__","forceClean":"TRUE","deviceTestMode":"noTest","deviceTestTimeSecs":""}}	id

Table 62 SDS Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			`deviceTestMode` – testOnly or noTest or testAndActivate. testOnly, Devices will be tested but not used. Later issue activateDevice to start using their capacity. `noTest`, Devices capacity will be used without any device testing.		
	/api/instances/Sds::{id}/action/setSdsName	POST	Required: `name`		
	/api/instances/Sds::{id}/action/setSdsPort	POST	Required: `sdsPort`		
	/api/instances/Sds::{id}/action/setDrlMode	POST	Required: `drlMode` - "Volatile" or "NonVolatile"		

Add an IP address to a SDS	/api/instances/Sds::{id}/action/addSdsIp	POST	Required:
			`ip`
			`role`
				When using the REST API with curl syntax, the command may not work as expected. Use curl syntax similar to the following example:	

Table 62 SDS Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
				curl -k -v -g --basic --user admin:$token -"Content-Type:application/json" -d '{"protectionDomainId":"ID", "sdsIpList":[{"SdsIp":{"ip":"IP","role":"all"}}] ,"deviceInfoList":[{"devicePath":"/dev/sdb","storagePoolId":"ID", "deviceName":"xyz"} >' https://GW-IP/api/types/Sds/instances	
Remove an IP address from a SDS	/api/instances/Sds::{id}/action/removeSdsIp	POST	Required: `ip`		
Set new role to existing IP address	/api/instances/Sds::{id}/action/setSdsIpRole	POST	Required: `sdsIpToSet` – IP that its role is about to change `newRole` – new role		
	/api/instances/Sds::{id}/action/startSdsNetworkTest	POST	Optional: numParallelMsgs, networkTestSizeInGb,	numParallelMsgs is a number in the range - 16. networkTestSizeInGb	

Table 62 SDS Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			networkTestLengthInSecs	value should be less than 10	
Clear error from all devices in a SDS	/api/instances/Sds::{id}/action/clearDevicesError	POST			
Activate all devices which completed their tests	/api/instances/Sds::{id}/action/activateDevices	POST			
Start devices test	/api/instances/Sds::{id}/action/startDevicesTest	POST	Optional: pattern - "writeSeq" or "writeRand" or "readSeq" or "readRand" or "last" ioSizeKB - a number between 4 and 512 specifying the single IO size in KB, totalIoSizeMB - a number between 1 and 32768 specifying the total IO size in MB, limitInSeconds	Default values: pattern - "readRand" ioSizeKB - 8 totalIoSizeMB - 128 limitInSeconds - 10 Performing device tests may fail application /.	
	/api/instances/ProtectionDomain::{id}/action/queryDevicesTestResults	POST			Response: List of DeviceTestResults
Limit the network	/api/instances/	POST	Required:	Overall limits size must be	setSdsNetworkLimits should be in

Table 62 SDS Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
bandwidth of all SDSs in the Protection Domain for various traffic types	ProtectionDomain::{id}/action/setSdsNetworkLimits		rebuildLimitInKbps AND\OR rebalanceLimitInKbp AND\OR overallLimitInKbps	larger than both rebuild and rebalance and all limits must be over 5MB	the ProtectionDomain action list
	/api/instances/ProtectionDomain::{id}/action/queryProtectionDomainNetwork	POST			Response: List of SdsConnectionState queryProtectionDomainNetwork should be in ProtectionDomain action list
	/api/instances/Sds::{id}/action/setSdsRmcacheEnabled	POST	Required: `rmcacheEnabled` – TRUE or FALSE		
	/api/instances/Sds::{id}/action/setSdsRmcacheSize	POST	Required: `rmcacheSizeInMB`		
	/api/instances/Sds::{id}/action/purgeRmcache	POST			
	/api/instances/Sds::{id}/action/setRmcacheFrozen	POST	Required: `rmcacheFrozen` - TRUE or FALSE		
	/api/instances/Sds::{id}/action/	POST			NetBwc list

Table 62 SDS Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	querySdsNetworkLatencyMeters				
	/api/instances/Sds::{id}/action/setNumOfIoBuffers	POST	Required: `numOfIoBuffers`		Deprecate (does not exist in version 2.0) The action setSdsPerformanceParameters can be used instead.
	/api/instances/Sds/action/queryDisconnectedSdss	POST			List of SDS objects which are disconnected, i.e. their MDM is disconnected or their membership state is not “joined”
	/api/instances/Sds::{id}/action/addSdsRfcacheDevice	POST	Required: `rfcacheDeviceInfoList` — list of: "devicePath" and "deviceName" Example: "rfcacheDeviceInfoList": [{"devicePath": "/dev/sdg","deviceName":"sdg"}, {"devicePath":"/dev/sdb","deviceName":"sdb"}]		
	/api/instances/Sds::{id}/action/clearRfcacheError	POST			
	/api/instances/	POST			

Table 62 SDS Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	Sds::{id}/action/enableSdsRfcache				
	/api/instances/Sds::{id}/action/disableSdsRfcache	POST			
	/api/instances/Sds::{id}/action/setSdsPerformanceParameters	POST	Required: `perfProfile` - "Default" or "HighPerformance"		

Example return output from `querySelectedStatistics`:

```
 Body: {"ids":["ead4729d00000000","ead4729e00000001"], "properties":
["rmcacheSkipCountLargeIo","primaryReadFromDevBwc","degradedHealthyV
acInKb","snapCapacityInUseInKb">
Response:
List of id and the requested statistics properties
{
"ead4729d00000000":{
"primaryReadFromDevBwc":{
"numOccured":0,
"totalWeightInKb":0,
"numSeconds":0
},
"rmcacheSkipCountLargeIo":0,
"snapCapacityInUseInKb":2097152,
"degradedHealthyVacInKb":0
},
"ead4729e00000001":{
"primaryReadFromDevBwc":{
"numOccured":0,
"totalWeightInKb":0,
"numSeconds":0
},
"rmcacheSkipCountLargeIo":0,
"snapCapacityInUseInKb":1048576,
"degradedHealthyVacInKb":0
}
}
```

Table 63 NetBwc

Properties	Type
sendSizeBwc	MosBwcQueryWindowSize

Table 63 NetBwc (continued)

Properties	Type
sendLatencyBwc	MosBwcQueryWindowLatency
sdsId	Long
sdsIp	String

Table 64 MosBwcQueryWindowSize

Properties	Type
numOccured	Long
totalWeightInKB	Long
numSeconds	Integer

Table 65 MosBwcQueryWindowLatency

Properties	Type
numOccured	Long
totalWeightInMicroseconds	Long
numSeconds	Integer

Storage Pool

Table 66 Storage Pool Object

Properties	Type	Required\Optional for POST
id	String	-
name	String	Optional
sparePercentage	Integer	Optional
rebuildEnabled	Boolean	Optional
rebalanceEnabled	Boolean	Optional
rebuildIoPriorityPolicy	unlimited or limitNumOfConcurrentIos or favorApplos or dynamicBwThrottling	-
rebalanceIoPriorityPolicy		-
rebuildIoPriorityNumOfConcurrentIosPerDevice	Integer	-
rebalanceIoPriorityNumOfConcurrentIosPerDevice	Integer	-

Table 66 Storage Pool Object (continued)

Properties	Type	Required\Optional for POST
rebuildIoPriorityBwLimitPerDeviceInKbps	Integer	-
rebalanceIoPriorityBwLimitPerDeviceInKbps	Integer	-
rebuildIoPriorityApplopsPerDeviceThreshold	Integer	-
rebalanceIoPriorityApplopsPerDeviceThreshold	Integer	-
rebuildIoPriorityAppBwPerDeviceThresholdInKbps	Integer	-
rebalanceIoPriorityAppBwPerDeviceThresholdInKbps	Integer	-
rebuildIoPriorityQuietPeriodInMsec	Integer	-
rebalanceIoPriorityQuietPeriodInMsec	Integer	-
numOfParallelRebuildRebalanceJobsPerDevice	Integer	-
protectionDomainId	String	Required
zeroPaddingEnabled	Boolean	Optional
useRmcache	Boolean	Optional
rmcacheWriteHandlingMode	Passthrough or Cached (default)	Optional
backgroundScannerMode	Disabled or DeviceOnly or DataComparison	-
backgroundScannerBWLimitKBps	Integer	-
checksumEnabled	Boolean	Required
useRfcache	Boolean	Required
useDAScache	Boolean	Required
capacityAlertHighThreshold	Integer	-
capacityAlertCriticalThreshold	Integer	-

Table 67 Storage Pool Parent

Name	Type	Note
ProtectionDomain	Protection Domain	

Table 68 Storage Pool Relationships

Name	Type	Note
Device	List of Device objects	
VTree	List of VTree objects	
Volume	List of Volume objects	
Statistics		

Storage Pool Statistics and Operations

Table 69 Storage Pool Statistics Object

Properties	Type	Note
numOfDevices	Long	
numOfVtrees	Long	
numOfVolumes	Long	
capacityLimitInKb	Long	
maxCapacityInKb	Long	
capacityInUseInKb	Long	
thickCapacityInUseInKb	Long	
thinCapacityInUseInKb	Long	
snapCapacityInUseInKb	Long	
snapCapacityInUseOccupiedInKb	Long	
unreachableUnusedCapacityInKb	Long	
protectedVacInKb	Long	
degradedHealthyVacInKb	Long	
degradedFailedVacInKb	Long	
failedVacInKb	Long	
inUseVacInKb	Long	
activeMovingInFwdRebuildJobs	Long	
pendingMovingInFwdRebuildJobs	Long	
activeMovingOutFwdRebuildJobs	Long	
pendingMovingOutFwdRebuildJobs	Long	
activeMovingInBckRebuildJobs	Long	

Table 69 Storage Pool Statistics Object (continued)

Properties	Type	Note
pendingMovingInBckRebuildJobs	Long	
activeMovingOutBckRebuildJobs	Long	
pendingMovingOutBckRebuildJobs	Long	
activeMovingInRebalanceJobs	Long	
pendingMovingInRebalanceJobs	Long	
activeMovingRebalanceJobs	Long	
pendingMovingRebalanceJobs	Long	
primaryVacInKb	Long	“primary” represents Master MDM in version 2.0 and higher
secondaryVacInKb	Long	“secondary” represents Slave MDM in version 2.0 and higher
primaryReadBwc	Bandwidth Counter	“primary” represents Master MDM in version 2.0 and higher
primaryReadFromDevBwc	Bandwidth Counter	“primary” represents Master MDM in version 2.0 and higher
primaryWriteBwc	Bandwidth Counter	“primary” represents Master MDM in version 2.0 and higher
secondaryReadBwc	Bandwidth Counter	“secondary” represents Slave MDM in version 2.0 and higher
secondaryReadFromDevBwc	Bandwidth Counter	“secondary” represents Slave MDM in version 2.0 and higher
secondaryWriteBwc	Bandwidth Counter	“secondary” represents Slave MDM in version 2.0 and higher
totalReadBwc	Bandwidth Counter	
totalWriteBwc	Bandwidth Counter	
fwdRebuildReadBwc	Bandwidth Counter	
fwdRebuildWriteBwc	Bandwidth Counter	
bckRebuildReadBwc	Bandwidth Counter	
bckRebuildWriteBwc	Bandwidth Counter	

Table 69 Storage Pool Statistics Object (continued)

Properties	Type	Note
rebalanceReadBwc	Bandwidth Counter	
rebalanceWriteBwc	Bandwidth Counter	
spareCapacityInKb	Long	
capacityAvailableForVolumeAllocationInKb	Long	
protectedCapacityInKb	Long	
degradedHealthyCapacityInKb	Long	
degradedFailedCapacityInKb	Long	
failedCapacityInKb	Long	
movingCapacityInKb	Long	
activeMovingCapacityInKb	Long	
pendingMovingCapacityInKb	Long	
fwdRebuildCapacityInKb	Long	
activeFwdRebuildCapacityInKb	Long	
pendingFwdRebuildCapacityInKb	Long	
bckRebuildCapacityInKb	Long	
activeBckRebuildCapacityInKb	Long	
pendingBckRebuildCapacityInKb	Long	
rebalanceCapacityInKb	Long	
activeRebalanceCapacityInKb	Long	
pendingRebalanceCapacityInKb	Long	
atRestCapacityInKb	Long	
numOfUnmappedVolumes	Long	
numOfMappedToAllVolumes	Long	
numOfThickBaseVolumes	Long	
numOfThinBaseVolumes	Long	
numOfSnapshots	Long	
numOfVolumesInDeletion	Long	
fixedReadErrorCount	Long	

Table 69 Storage Pool Statistics Object (continued)

Properties	Type	Note
BackgroundScanCompareCount	Long	
BackgroundScannedInMB	Long	
thinCapacityAllocatedInKm	Long	
rmPendingAllocatedInKb	Long	
semiProtectedVacInKb	Long	
inMaintenanceVacInKb	Long	
activeMovingInNormRebuildJobs	Long	
activeMovingOutNormRebuildJobs	Long	
pendingMovingInNormRebuildJobs	Long	
pendingMovingOutNormrebuildJobs	Long	
primaryReadFromRmcacheBwc	Bandwidth Counter	“primary” represents Master MDM in version 2.0 and higher
secondaryReadFromRmcacheBwc	Bandwidth Counter	“secondary” represents Slave MDM in version 2.0 and higher
normRebuildReadBwc	Bandwidth Counter	
normRebuildWriteBwc	Bandwidth Counter	
rfcacheReadsReceived	Long	
rfcacheWritesReceived	Long	
rfcacheAvgReadTime	Long	
rfcacheAvgWriteTime	Long	
rfcacheSourceDeviceReads	Long	
rfcacheSourceDeviceWrites	Long	
rfacheReadHit	Long	
rfcacheReadMiss	Long	
rfcacheWriteMiss	Long	
rfcacheIosSkipped	Long	
rfcacheReadsSkipped	Long	
rfcacheReadsSkippedAlignedSizeTooLarge	Long	
rfcacheReadsSkippedMaxIoSize	Long	

Table 69 Storage Pool Statistics Object (continued)

Properties	Type	Note
rfcacheReadsSkippedHeavyLoad	Long	
rfcacheReadsSkippedStucklo	Long	
rfcacheReadsSkippedLowResources	Long	
rfcacheReadsSkippedInternalError	Long	
rfcacheReadsSkippedLocklos	Long	
rfcacheWritesSkippedMaxloSize	Long	
rfcacheWritesSkippedHeavyLoad	Long	
rfcacheWritesSkippedStucklo	Long	
rfcacheWritesSkippedLowResource s	Long	
rfcacheWritesSkippedInternalError	Long	
rfcacheWritesSkippedCacheMiss	Long	
rfcacheSkippedUnlinedWrite	Long	
rfcacheIoErrors	Long	
rfcacheReadsFromCache	Long	
rfcachelosOutstanding	Long	
rfcacheReadsPending	Long	
rfcacheWritePending	Long	
semiProtectedCapacityInKb	Long	
inMaintenanceCapacityInKb	Long	
normRebuildCapacityInKb	Long	
activeNormRebuildCapacityInKb	Long	
pendingNormRebuildCapacityInKb	Long	

Table 70 Storage Pool Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/StoragePool/instances	GET			List of all StoragePool objects
	/api/instances/StoragePool::{id}	GET			StoragePool object

Table 70 Storage Pool Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/StoragePool::{id}/relationships/Device	GET			List of Device objects that are connected to the Storage Pool
	/api/instances/StoragePool::{id}/relationships/VTree	GET			List of VTree objects that are connected to the Storage Pool
	/api/instances/StoragePool::{id}/relationships/Volume	GET			List of Volume objects that are connected to the Storage Pool
	/api/instances/StoragePool::{id}/relationships/Statistics	GET			List of Statistics for the Storage Pool

	/api/types/StoragePool/instances/action/querySelectedStatistics/	POST	Required:
			`ids` - list of objects IDs
			or
			`allIds` - with empty value
			`properties` - list of properties to fetch (to see available properties, see Statistics)
				If `allIds` appears, ids list is ignored and a list of selected properties is returned for every StoragePool.	List of id and the selected properties
					List of:
					id and the selected properties
					{"ids": ["a7e798d100000000"], "properties": ["thickCapacityInUseInKb","thinCapacityInUseInKb" >
					Response:
					{
					""a7e798d100000000":{

Table 70 Storage Pool Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
					o "thickCapacityInUseInKb":1, o "thinCapacityInUseInKb":0 } }

	/api/types/StoragePool/instances/action/queryIdByKey	POST	Required:
			`name`
			AND
			`protectionDomainId` or `protectionDomainName`
					id

	/api/types/StoragePool/instances/action/queryBySelectedIds	POST	Required: `ids`		List of StoragePool objects matching the ids
Remove a StoragePool	/api/instances/StoragePool::{id}/action/removeStoragePool	POST			
	/api/types/StoragePool/instances	POST	According to "Required \Optional for POST" column in StoragePool object. Zero padding must be configured before devices are added.		id
	/api/instances/StoragePool::{id}/action/setStoragePoolName	POST	Required: `name`		

Table 70 Storage Pool Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/StoragePool::{id}/action/setSparePercentage	POST	Required: `sparePercentage`		
	/api/instances/StoragePool::{id}/action/queryDevicesTestResults	POST			Response: List of DeviceTestResults
	/api/instances/StoragePool::{id}/action/setZeroPaddingPolicy	POST	Required: `zeroPadEnabled` – TRUE or FALSE	Zero padding must be configured before devices are added.	
	/api/instances/StoragePool::{id}/action/setUseRmcache	POST	Required: `useRmcache` – `TRUE or FALSE (Default)`		
	/api/instances/StoragePool::{id}/action/setRebuildIoPriorityPolicy	POST	Required: `policy` - unlimited or limitNumOfConcurrentIos or favorApplos or dynamicBwThrottling Optional: `numOfConcurrentIosPerDevice` - The maximum number of concurrent rebuild IOs per device. Valid range: 1–20, `bwLimitPerDeviceInKbp` - The maximum bandwidth of		

Table 70 Storage Pool Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			rebuild IOs, in KB per second, per device. Valid range: 1024–1048576, `appIopsPerDeviceThreshold` - The application IOPs threshold above which rebuild IO throttling will be applied. Valid range: 1–1024, `appBwPerDeviceThresholdInKbps` - The application IO bandwidth threshold, in KB per second,		
			above which rebuild IO throttling will be applied. Valid range: 1-102400, `quietPeriodInMsec` - If the application IOs are below the IOPs and bandwidth threshold in this period, given in milliseconds, no throttling will be applied to the rebuild IOs. Valid range: 10-60000		
	/api/instances/StoragePool::{id}/action/setRebalanceIoPriorityPolicy	POST	Required: `policy` - unlimited or limitNumOfConcurrentIos or favorAppIos or dynamicBwThrottling		

Table 70 Storage Pool Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			Optional: `numOfConcurrentIosPerDevice` - The maximum number of concurrent rebuild IOs per device. Valid range: 1–20, `bwLimitPerDeviceInKbp` - The maximum bandwidth of rebuild IOs, in KB per second, per device. Valid range: 1024–1048576, `appIopsPerDeviceThreshold` - The application IOPs threshold above which rebuild IO throttling will be applied. Valid range: 1-1024, `appBwPerDeviceThresholdInKbps` - The application IO bandwidth threshold, in KB per second, above which rebuild IO throttling will be		
			applied. Valid range: 1–102400, `quietPeriodInMsec` - If the application IOs are below the IOPs and bandwidth threshold in this		

Table 70 Storage Pool Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			period, given in mgiven in milliseconds, no throttling will be applied to the rebuild IOs. Valid range: 10–60000		
	/api/instances/StoragePool::{id}/action/setRmcacheWriteHandlingMode	POST	`rmcacheWriteHandlingMode` - Passthrough or Cached		
	/api/instances/StoragePool::{id}/action/setRebuildEnabled	POST	Required: `rebuildEnabled—TRUE or FALSE`		
	/api/instances/StoragePool::{id}/action/setRebalanceEnabled	POST	Required: `rebalanceEnabled—TRUE or FALSE`		
	/api/instances/StoragePool::{id}/action/setRebuildRebalanceParallelism	POST	Required: `limit—The maximum number of concurrent rebuild and rebalance activities. Number between 1-10.`		
	/api/instances/StoragePool::{id}/action/enableBackgroundDeviceScanner	POST	Required: `scannerMode` - "DeviceOnly" or "DataComparison" Optional:	DeviceOnly - Perform read operations. Fix from peer on errors. DataComparison - Read and compare	

Table 70 Storage Pool Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			`bandwidthLimitInKbpsPerDevice`	data content with peer.	
	/api/instances/StoragePool::{id}/action/disableBackgroundDeviceScanner	POST			
	/api/instances/StoragePool::{id}/action/resetBackgroundDeviceScannerErrorCounters	POST	Optional: `dataCompareErrorCounter` - TRUE or FALSE `correctedReadErrorCounter` - TRUE or FALSE	Default is FALSE for both counters	
	/api/instances/StoragePool::{id}/action/setChecksumEnabled	POST	Required: `checksumEnabled` - TRUE or FALSE		
	/api/instances/StoragePool::{id}/action/setRfcacheUsage	POST	Required: `useRfcache` - TRUE or FALSE		
	/api/instances/StoragePool::{id}/action/setCapacityAlertThresholds	POST	Required: capacityAlertHighThresholdPercent AND\OR capacityAlertCriticalThresholdPercent The high threshold value must be lower than the critical threshold value.		

Table 70 Storage Pool Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/StoragePool::{id}/action/resetDeviceOscillatingCounter	POST	Required: `allCounters` - TRUE or FALSE OR `failureCounter` - one of the following: DevLongSuccessfullos		
	/api/instances/StoragePool::{id}/action/setDeviceOscillatingCounterParameters	POST	Required: `failureCounter` - see previous operation for possible values `windowType` - short or medium or long `windowInterval` - number between 1 and 86400 failuresThreshold		

Device

Table 71 Device Object

Properties	Type	Required\Optional for POST
id	String	-
name	String	Optional
deviceCurrentPathname	String	Required
deviceOriginalPathname	String	-
deviceState	InitialTest or InitialTestDone or Normal or NormalTesting or RemovePending	-
errorState	None or Error	-

Table 71 Device Object (continued)

Properties	Type	Required\Optional for POST
capacityLimitInKb	Long	Optional
maxCapacityInKb	Long	-
storagePoolId	String	Required
longSuccessfulIos	OscillatingCounter	Optional
rfcacheErrorDeviceDoesNotExist	Boolean	Optional
sdsId	String	Required
updateConfiguration	Boolean	Optional

Table 72 Device Parent

Name	Type	Note
StoragePool	StoragePool	
SDS	SDS	

Table 73 Device Relationships

Name	Type	Note
Statistics		

Table 74 OscillatingCounter

Name	Type	Note
shortWindow	OscillatingCounterWindow	
mediumWindow	OscillatingCounterWindow	
longWindow	OscillatingCounterWindow	

Table 75 OscillatingCounterWindow

Name	Type	Note
threshold	Integer	
windowSizeInSec	Integer	
lastOscillationCount	Integer	
lastOscillationiTime	Long	
maxFailuresCount	Integer	

Table 75 OscillatingCounterWindow (continued)

Name	Type	Note
maxFailuresTime	Long	

Device Statistics and Operations

Note

If the background device scanner is enabled, read statistics are dramatically affected.

Table 76 Device Statistics Object

Properties	Type	Note
fixedReadErrorCount	Long	
avgReadSizeInBytes	Long	
avgWriteSizeInBytes	Long	
avgReadLatencyInMicrosec	Long	
avgWriteLatencyInMicrosec	Long	
capacityInUseInKb	Long	
thickCapacityInUseInKb	Long	
thinCapacityInUseInKb	Long	
snapCapacityInUseInKb	Long	
snapCapacityInUseOccupiedInKb	Long	
unreachableUnusedCapacityInKb	Long	
protectedVacInKb	Long	
degradedHealthyVacInKb	Long	
degradedFailedVacInKb	Long	
failedVacInKb	Long	
inUseVacInKb	Long	
activeMovingInFwdRebuildJobs	Long	
pendingMovingInFwdRebuildJobs	Long	
activeMovingOutFwdRebuildJobs	Long	
pendingMovingOutFwdRebuildJobs	Long	
activeMovingInBckRebuildJobs	Long	
pendingMovingInBckRebuildJobs	Long	

Table 76 Device Statistics Object (continued)

Properties	Type	Note
activeMovingOutBckRebuildJobs	Long	
pendingMovingOutBckRebuildJobs	Long	
activeMovingInRebalanceJobs	Long	
pendingMovingInRebalanceJobs	Long	
activeMovingRebalanceJobs	Long	
pendingMovingRebalanceJobs	Long	
primaryVacInKb	Long	“primary” represents the Master MDM in version 2.0 and higher
secondaryVacInKb	Long	“secondary” represents the Slave MDM in version 2.0 and higher
primaryReadBwc	Bandwidth Counter	“primary” represents the Master MDM in version 2.0 and higher
primaryReadFromDevBwc	Bandwidth Counter	“primary” represents the Master MDM in version 2.0 and higher
primaryWriteBwc	Bandwidth Counter	“primary” represents the Master MDM in version 2.0 and higher
secondaryReadBwc	Bandwidth Counter	“secondary” represents the Slave MDM in version 2.0 and higher
secondaryReadFromDevBwc	Bandwidth Counter	“secondary” represents the Slave MDM in version 2.0 and higher
secondaryWriteBwc	Bandwidth Counter	“secondary” represents the Slave MDM in version 2.0 and higher
totalReadBwc	Bandwidth Counter	
totalWriteBwc	Bandwidth Counter	
fwdRebuildReadBwc	Bandwidth Counter	
fwdRebuildWriteBwc	Bandwidth Counter	
bckRebuildReadBwc	Bandwidth Counter	
bckRebuildWriteBwc	Bandwidth Counter	

Table 76 Device Statistics Object (continued)

Properties	Type	Note
rebalanceReadBwc	Bandwidth Counter	
rebalanceWriteBwc	Bandwidth Counter	
BackgroundScanCompareCount	Long	
BackgroundScannedInMB	Long	
thinCapacityAllocatedInKm	Long	
rmPendingAllocatedInKb	Long	
semiProtectedVacInKb	Long	
inMaintenanceVacInKb	Long	
activeMovingInNormRebuildJobs	Long	
activeMovingOutNormRebuildJobs	Long	
pendingMovingInNormRebuildJobs	Long	
pendingMovingOutNormrebuildJobs	Long	
primaryReadFromRmcacheBwc	Bandwidth Counter	“primary” represents the Master MDM in version 2.0 and higher
secondaryReadFromRmcacheBwc	Bandwidth Counter	“secondary” represents the Slave MDM in version 2.0 and higher
normRebuildReadBwc	Bandwidth Counter	
normRebuildWriteBwc	Bandwidth Counter	
rfcacheReadsReceived	Long	
rfcacheWritesReceived	Long	
rfcacheAvgReadTime	Long	
rfcacheAvgWriteTime	Long	
rfcacheSourceDeviceReads	Long	
rfcacheSourceDeviceWrites	Long	
rfacheReadHit	Long	
rfcacheReadMiss	Long	
rfcacheWriteMiss	Long	
rfcacheIosSkipped	Long	
rfcacheReadsSkipped	Long	

Table 76 Device Statistics Object (continued)

Properties	Type	Note
rfcacheReadsSkippedAlignedSizeTooLarge	Long	
rfcacheReadsSkippedMaxIoSize	Long	
rfcacheReadsSkippedHeavyLoad	Long	
rfcacheReadsSkippedStuckIo	Long	
rfcacheReadsSkippedLowResources	Long	
rfcacheReadsSkippedInternalError	Long	
rfcacheReadsSkippedLockIos	Long	
rfcacheWritesSkippedMaxIoSize	Long	
rfcacheWritesSkippedHeavyLoad	Long	
rfcacheWritesSkippedStuckIo	Long	
rfcacheWritesSkippedLowResources	Long	
rfcacheWritesSkippedInternalError	Long	
rfcacheWritesSkippedCacheMiss	Long	
rfcacheSkippedUnlinedWrite	Long	
rfcacheIoErrors	Long	
rfcacheReadsFromCache	Long	
rfcacheIosOutstanding	Long	
rfcacheReadsPending	Long	
rfcacheWritePending	Long	

Table 77 Device Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/Device/instances	GET		If zero padding is required , you must enable it before adding any devices to the Storage Pool.	List of all Deviceobjects

Table 77 Device Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/Device::{id}	GET			Deviceobject
	/api/instances/Device::{id}/relationships/Statistics	GET			List of Statistics for the Device

	/api/types/Device/instances/action/querySelectedStatistics	POST	Required:
			`ids`- list of object IDs
			or
			`allIds`- with empty value
			`properties`- list of properties to fetch (to see available properties, see Statistics)
				If `allIds`appears, ids list is ignored and a list of selected properties is returned for every Device.	List of id and the selected properties
					{"ids": ["a7e798d100000000"], "properties": ["avgReadLatencyInMicrosec","avgWriteLatencyInMicrosec">
					Response:
					{
					""a7e798d100000000":{
					o "avgReadLatencyInMicrosec":1,
					o "avgWriteLatencyInMicrosec":0
					}
					}

	/api/types/Device/instances/action/queryIdByKey	POST	Required:
			`name`
			AND
			`sdsId`or `sdsName`
					id

	/api/types/Device/instances/action/queryBySelectedIds	POST	Required: `ids`		Matching ids

Table 77 Device Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
Remove a device	/api/instances/Device::{id}/action/removeDevice	POST	Required: `force`- TRUE or FALSE.		
	/api/instances/Device::{id}/action/abortRemoveDevice	POST		This request attempts to abort the removal of a device. If done at a late stage, this request might fail and the device will be removed.	
	/api/types/Device/instances	POST	According to "Required \Optional for POST" column in Device object. Plus: Optional: `testTimeSecs` - The maximal test run-time in seconds or 128 MB data. Default limit is 10 seconds. Not relevant for noTest. `testMode`– testOnly or noTest or testAndActivate. testOnly, Devices will be tested but not used. Later, issue activateDevice to start using their capacity.	If testMode is not part of the request body, a read and write test will be run on devices before their capacity will be used.	id

Table 77 Device Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			noTest, Devices capacity will be used without any device testing.		
	/api/instances/Device::{id}/action/setDeviceCapacityLimit	POST	Required: `capacityLimitInGB`		
	/api/instances/Device::{id}/action/clearDeviceError	POST		Clear device error state without checking if the device error was fixed. If the error continues to exist, the device will return to an error state as soon as it is accessed.	
Activate a device which completed its test	/api/instances/Device::{id}/action/activateDevice	POST	Required: `sdsId`	Device created with testOnly parameter will not be used by the system until activateDevice is sent.	
	/api/instances/Device::{id}/action/queryDeviceTestResults	POST			Response: List of DeviceTestResults
	/api/instances/Device::{id}/action/setDeviceName	POST	Required: `newName`		

Table 77 Device Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/Device::{id}/action/updateDeviceOriginalPathname	POST		Sets deviceOriginalPathname to the value of deviceCurrentPathname	
	/api/instances/Device/action/queryFailedDevices	POST			List of Device objects which are in error state.

Volume

Table 78 Volume Object

Properties	Type	Required\Optional for POST
id	String	-
name	String	Optional
volumeSizeInKb	Long	Required
isObfuscated	Deprecated in v2.0	Deprecated in v2.0
creationTime	Long	-
volumeType	ThickProvisioned or ThinProvisioned or Snapshot	Optional (only ThickProvisioned and ThinProvisioned)
consistencyGroupId	String	-
mappingToAllSdcsEnabled	Boolean	-
mappedSdcInfoList	List of SdcMappingInfo	-
ancestorVolumeId	String	-
vtreeId	String	-
storagePoolId	String	Required
useRmcache	Boolean	Optional

Table 79 Volume Parent

Name	Type	Note
VTree	VTree	
StoragePool	StoragePool	
Volume	Volume	Ancestor volume

Table 80 Volume Relationships

Name	Type	Note
Statistics		

Volume Statistics and Operations

Table 81 Volume Statistics Object

Properties	Type	Note
numOfMappedSdcs	Long	
numOfChildVolumes	Long	
numOfDescendantVolumes	Long	
userDataReadBwc	Bwc	
userDataWriteBwc	Bwc	

Table 82 SdcMappingInfo

Properties	Type	Note
sdcId	String	
sdcIp	String	
limitIops	Long	
limitBwInMbps	Long	

Table 83 Volume Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/Volume/instances	GET			List of all Volume objects
	/api/instances/Volume::{id}	GET			Volume object

Table 83 Volume Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/Volume::{id}/relationships/Statistics	GET			List of Statistics for the Volume

	/api/types/Volume/instances/action/querySelectedStatistics	POST	Required:
			`ids` - list of object IDs
			or
			`allIds` - with empty value
			`properties` - list of properties to fetch (to see available properties, see Statistics)
				If `allIds` appears, ids list is ignored and a list of selected properties is returned for every Volume.	List of id and the selected properties
					properties
					For example:
					Body:
					{"ids": ["a7e798d100000000"], "properties": ["numOfMappedSdcs","numOfChildVolumes">
					Response:
					{
					""a7e798d100000000":{
					o "numOfMappedSdcs":1,
					o "numOfChildVolumes":0
					}
					}

	/api/types/Volume/instances/action/queryIdByKey	POST	Required: `name`		
	/api/types/Volume/instances/action/queryBySelectedIds	POST	Required: `ids`		
	/api/instances/	POST	Required:		

Table 83 Volume Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	Volume::{id}/action/removeVolume		`removeMode` – "ONLY_ME" or "INCLUDING_DESCENDANTS" or "DESCENDANTS_ONLY" or "WHOLE_VTREE"		
	/api/types/Volume/instances	POST	According to “Required \Optional for POST” column in Volume object.		id
	/api/instances/Volume::{id}/action/setVolumeSize	POST	Required: `sizeInGB`	New volume size in GB. Basic allocation granularity is 8 GB	
	/api/instances/Volume::{id}/action/setVolumeName	POST	Required: `newName`		
Map a volume to one or all SDC nodes	/api/instances/Volume::{id}/action/addMappedSdc	POST	Required: `sdcId` – SDC id or `guid` – SDC guid Optional: `allowMultipleMappings` - "TRUE" or "FALSE"	Map - exposes a volume to all SDC nodes, the exposure will also include SDC nodes that will be added after the execution of this request.	
Unmap a volume from	/api/instances/	POST	`Required:`	If allSdcs appears,	

Table 83 Volume Operations (continued)

Operation	URI	Method	Parameters	Notes	Return

one or all SDC nodes	Volume::{id}/action/removeMappedSdc		`sdcId - SDC id`
			`or`
			`guid` - SDC guid
			or
			`allSdcs with empty value`
			`Optional:`
			`ignoreScsiInitiators - "TRUE" or "FALSE"`
			`skipApplianceValidation (TRUE\FALSE)`
				unmap volume from all SDCs.	

Set limits to the IOPS and bandwidth that one SDC generates for the specified volume.	/api/instances/Volume::{id}/action/setMappedSdcLimits	POST	Required: `sdcId` or `guid` – SDC GUID Optional: `bandwidthLimitInKbps` - Limit the volume network bandwidth. A positive number in granularity of 1024 Kbps. 0 is unlimited `iopsLimit` - Limit the volume IOPS. The number of IOPS must be larger than 10. 0 is unlimited.		
	/api/instances/Volume::{id}/action/	POST	Required: `useRmcache` - TRUE or FALSE (default)		

Table 83 Volume Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	setVolumeUseRmcache				
Revert snapshot	/api/instances/Volume::{id}/action/revertSnapshot	POST	Required: `volumeId` - Volume from URL and volumeId must be from the same VTree. `revertMode` - "move" or "copy" or "toggle"	revertMode: "move" - Volume from URL gets the content of the volumeId; volumeId is no longer available. "copy" - Volume from URL gets the content of the volumeId; volumeId remains untouched. "toggle" - the content of the Volume from URL and the volumeId are switched.	

VTree

Table 84 VTree Object

Properties	Type	Note
id	String	
name	String	
baseVolumeId	String	
storagePoolId	String	

Table 85 VTree Parent

Name	Type	Note
Volume	Volume	Base volume

Table 85 VTree Parent (continued)

Name	Type	Note
StoragePool	Storage Pool on page 471	

Table 86 VTree Relationships

Name	Type	Note
Volume	List of Volume objects	
Statistics		

VTree Statistics and Operations

Table 87 VTree Statistics Object

Properties	Type	Note
numOfVolumes	Long	
netCapacityInUseInKb	Long	
baseNetCapacityInUseInKb	Long	
snapNetCapacityInUseInKb	Long	
trimmedCapacityInKb	Long	

Example of return output for `querySelectedStatistics`:

```
 Body: {"ids":["a7e798d100000000"], "properties":
["netCapacityInUseInKb","numOfVolumes">
Response:
{
"a7e798d100000000":{
"numOfVolumes":1,
"netCapacityInUseInKb":0
}
}
```

SDC

Table 88 SDC Object

Properties	Type	Required\Optional for POST
id	String	-
name	String	Optional
sdclp	String	Required

Table 88 SDC Object (continued)

Properties	Type	Required\Optional for POST
sdcGuid	String	-
onVmWare	Boolean	-
mdmConnectionState	Connected or Disconnected	-
systemId	String	-
sdcApproved	Boolean	-
versionInfo	String	-
perfParams	SdcPerfParams	-
currProfilePerfParams	SdcPerfParams	-
memoryAllocationFailure	OscillatingCounter	-
socketAllocationFailure	OscillatingCounter	-

Table 89 SDC Parent

Name	Type	Note
System	System	

Table 90 SDC Relationships

Name	Type	Note
Volume	List of Volume objects	Mapped volumes
Statistics		

Table 91 SdcPerfParams

Properties	Type	Note
sdcTcpSendBufferSize	Integer	
sdcTcpReceiveBufferSize	Integer	
sdcNumberSocketsPerSdsIp	Integer	
sdcNumberNetworkOsThreads	Integer	
sdcMaxInflightRequests	Integer	
sdcMaxInFlightData	Integer	
sdcNumberIoRetries	Integer	
sdcVolumeStatisticsInterval	Integer	

Table 91 SdcPerfParams (continued)

Properties	Type	Note
sdcOptimizeZeroBuffers	Boolean	
sdcNumberUnmapBlocks	Integer	
sdcNumberNonIoOsThreads	Integer	
perfProfile	Default or HighPerformance	

SDC Statistics and Operations

Table 92 SDC Statistics Object

Properties	Type	Note
numOfMappedVolumes	Long	
userDataReadBwc	Bandwidth Counter	
userDataWriteBwc	Bandwidth Counter	

Table 93 SDC Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/Sdc/instances	GET			List of all SDC objects **Note** Only one IP address is returned for each SDC in the sdcIp field, even if it has more than one IP address. The address that is returned is the one that is used to communicate with the MDM.
	/api/instances/Sdc::{id}	GET			SDC object
	/api/instances/Sdc::{id}/	GET			List of Volume objects that are connected to the SDC

Table 93 SDC Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	relationships/Volume				
	/api/instances/Sdc::{id}/relationships/Statistics	GET			List of Statistics for the SDC

	/api/types/Sdc/instances/action/querySelectedStatistics	POST	Required:
			`ids` - list of object IDs
			or
			`allIds` - with empty value
			`properties` - list of properties to fetch (to see available properties see Statistics)
				If `allIds` appears, ids list is ignored and a list of selected properties is returned for every SDC.	List of id and the selected properties
					For example:
					Body:
					{"ids": ["a7e798d100000000"], "properties": ["numOfMappedVolumes">
					Response:
					{
					""a7e798d100000000":{
					o "numOfMappedVolumes":1}
					}

	/api/types/Sdc/instances/action/queryAllApprovedSdc	POST			List of approved SDC objects
	/api/types/Sdc/instances/action/queryIdByKey	POST	Required: `IP address`		id
	/api/types/Sdc/instances/action/queryBySelectedIds	POST	Required: `ids`		List of SDC objects matching the ids

Table 93 SDC Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/Sdc::{id}/action/removeSdc	POST			
	/api/instances/Sdc::{id}/action/setSdcName	POST	Required: `sdcName`		
	/api/instances/Sdc::{id}/action/queryOscillatingNetworkCounters	POST			
	/api/types/Sdc/instances	POST	According to "Required \Optional for POST" column in SDC object.		
	/api/instances/Sdc::{id}/action/setSdcPerformanceParameters	POST	Required: `perfProfile` - "Default" or "HighPerformance"		

User Objects

Table 94 User Object

Properties	Type	Required\Optional for POST
id	String	-
name	String	Required
userRole	Monitor, Configure, Administrator, Superuser, Security, FrontendConfig or BackendConfig	Required
passwordChangeRequired	Boolean	-

Table 94 User Object (continued)

Properties	Type	Required\Optional for POST
systemId	String	-

Table 95 User Parent

Name	Type	Note
System	System	

Table 96 User Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/User/instances	GET			List of all User objects
	/api/instances/User::{id}	GET			User object
	/api/instances/User::{id}/action/removeUser	POST			
	/api/types/User/instances	POST	According to "Required \Optional for POST" column in User object.	User cannot be created with the userRole Superuser	id defaultPassword
	/api/instances/User::{id}/action/setUserRole	POST	Required: `userRole`		
	/api/instances/User/action/setPassword	POST	Required: `oldPassword`, `newPassword`		
	/api/instances/User::{id}/action/resetPassword				defaultPassword

Fault Set

Table 97 Fault Set Object

Properties	Type	Required\Optional for POST
id	String	-
name	String	Optional
protectionDomainId	String	Required

Table 98 Fault Set Parent

Name	Type	Note
ProtectionDomain	Protection Domain on page 438	

Table 99 Fault Set Relationships

Name	Type	Note
SDS	List of SDS objects	
Statistics		

Fault Set Statistics and Operations

Table 100 Fault Set Statistics Object

Properties	Type	Note
numOfSds	Long	
rfcacheFdReadsReceived	Long	
rfcacheFdWritesReceived	Long	
rfcacheFdInlightReads	Long	
rfcacheFdInlightWrites	Long	
rfcacheFdReadTimeGreater500Millis	Long	
rfcacheFdReadTimeGreater1Sec	Long	
rfcacheFdReadTimeGreater5Sec	Long	
rfcacheFdReadTimeGreater1Min	Long	
rfcacheFdWriteTimeGreater500Millis	Long	
rfcacheFdWriteTimeGreater1Sec	Long	

Table 100 Fault Set Statistics Object (continued)

Properties	Type	Note
rfcacheFdWriteTimeGreater5Sec	Long	
rfcacheFdWriteTimeGreater1Min	Long	
rfcacheFdAvgReadTime	Long	
rfcacheFdAvgWriteTime	Long	
rfcacheFdIoErrors	Long	
rfcacheFdCacheOverloaded	Long	
rfcacheFdMonitorErrorStucklo	Long	
rfcachePoolNumSrcDevs	Long	
rfcachePoolNumCacheDevs	Long	
rfcachePoolSize	Long	
rfcachePoolReadHit	Long	
rfcachePoolReadMiss	Long	
rfcachePoolWriteHit	Long	
rfcachePoolWriteMiss	Long	
rfcachePoolCachePages	Long	
rfcachePoolPagesInuse	Long	
rfcachePoolEvictions	Long	
rfcachePoolInLowMemoryCondition	Long	
rfcachePoolIoTimeGreater1Min	Long	
rfcachePoolLockTimeGreater1Sec	Long	
rfcachePoolSuspendedIos	Long	
rfcachePoolLowResourcesInitiatedPassthroughMode	Long	
rfcachePoolosOutstanding	Long	
rfcachePoolReadsPending	Long	
rfcachePoolWritePending	Long	
rfcachePoolSuspendedPequestsRedundantSearchs	Long	
rfcachePoolReadPendingG1Sec	Long	
rfcachePoolReadPendingG10Millis	Long	
rfcachePoolReadPendingG1Millis	Long	

Table 100 Fault Set Statistics Object (continued)

Properties	Type	Note
rfcachePoolReadPendingG500Micro	Long	
rfcachePoolWritePendingG1Sec	Long	
rfcachePoolWritePendingG10Millis	Long	
rfcachePoolWritePendingG1Millis	Long	
rfcachePoolWritePendingG500Micro	Long	
rfcachePoolSourceIdMismatch	Long	

Table 101 Fault Set Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/FaultSet/instances	GET			List of all Fault Set on page 506 objects
	/api/instances/FaultSet::{id}	GET			Fault Set on page 506 object
	/api/instances/FaultSet::{id}/relationships/Sds	GET			List of SDS objects that are connected to the Fault Set
	/api/instances/FaultSet::{id}/relationships/Statistics	GET			List of Statistics for the Fault Set
	/api/types/FaultSet/instances/action/querySelectedStatistics	POST	Required: ids - list of object IDs or allIds - with empty value. properties - list of properties to fetch (to see available properties, see Statistics)	If allIds appears, ids list is ignored and a list of selected properties is returned for every FaultSet	List of id and the selected properties For example: Body: {"ids": ["a7e798d100000000"], "properties":["rfcachePoolReadHit","numOfVolumes"> Response: {

Table 101 Fault Set Operations (continued)

Operation	URI	Method	Parameters	Notes	Return

					""a7e798d100000000":{
					o "rfcachePoolReadHit":1,
					o "rfcachePoolEvictions":0
					}
					}

	/api/types/FaultSet/instances/action/queryIdByKey	POST	Required: name AND protectionDomainId or protectionDomainName		id
	/api/types/FaultSet/instances/action/queryBySelectedIds	POST	Required: ids		List of Fault Set on page 506 objects matching the ids
	/api/instances/FaultSet::{id}/action/removeFaultSet	POST			
	/api/types/FaultSet/instances	POST	According to “Required \Optional for POST” column in Fault Set on page 506 object.		id
	/api/instances/FaultSet::{id}/action/clearFaultSet	POST			
	/api/instances/	POST	Required: newName		

Table 101 Fault Set Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	FaultSet::{id}/action/setFaultSetName				
	/api/instances/FaultSet::{id}/action/setSdsPerformanceParameters	POST	Required: perfProfile - "Default" or "HighPerformance"		

RFcache Device

Table 102 RFcache Device Objects

Properties	Type	Required/Optional for POST
id	String	-
name	String	Required
sdsId	Long	Required
serviceCurrentPathname	String	-
serviceOriginalPathname	String	-
rfcacheErrorDeviceDoesNotExist	Boolean	-
rfcacheErrorStackl	Boolean	-
rfcacheErrorHeavyLoadCacheSkip	Boolean	-
rfcacheErrorCardIoError	Boolean	-

Table 103 RFcache Device Parent

Name	Type	Note
Sds	SDS	

rfcacheDevice Statistics and Operations

Table 104 rfcacheDevice Statistics Object

Properties	Type	Note
rfcacheFdReadsReceived	Long	

Table 104 rfcacheDevice Statistics Object (continued)

Properties	Type	Note
rfcacheFdWritesReceived	Long	
rfcacheFdInlightReads	Long	
rfcacheFdInlightWrites	Long	
rfcacheFdReadTimeGreater500Millis	Long	
rfcacheFdReadTimeGreater1Sec	Long	
rfcacheFdReadTimeGreater5Sec	Long	
rfcacheFdReadTimeGreater1Min	Long	
rfcacheFdWriteTimeGreater500Millis	Long	
rfcacheFdWriteTimeGreater1Sec	Long	
rfcacheFdWriteTimeGreater5Sec	Long	
rfcacheFdWriteTimeGreater1Min	Long	
rfcacheFdAvgReadTime	Long	
rfcacheFdAvgWriteTime	Long	
rfcacheFdIoErrors	Long	
rfcacheFdCacheOverloaded	Long	
rfcacheFdMonitorErrorStucklo	Long	

Table 105 rfcacheDevice Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/instances/RfcacheDevice::{id}/action/removeSdsRfcacheDevice	POST			
	/api/instances/RfcacheDevice::{id}/action/updateSdsRfcacheDeviceOriginalPath	POST			
	/api/instances/RfcacheDevice::{id}/action/	POST	Required: newName		

Table 105 rfcacheDevice Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
	renameSdsRfcacheDevice				
	/api/types/RfcacheDevice /instances	GET			List of all RfcacheDevice objects
	/api/instances/RfcacheDevice::{id}	GET			RfcacheDevice object

	/api/types/RfcacheDevice/instances/action/querySelectedStatistics	POST	Required:
			ids - list of objects ids
			or
			allIds - with empty value.
			- list of properties to fetch (to see available properties see Statistics)
				If allIds appears, ids list is ignored and a list of selected properties is returned for every Device.	List of:
					id and the selected properties
					`{"ids": ["a7e798d100000000"], "properties": ["rfcacheFdWritesReceived","rfcacheFdIoErrors">`
					Response:
					`{`
					`"a7e798d100000000":{`
					`"rfcacheFdWritesReceived": 1,`
					`"rfcacheFdIoErrors":0`
					`}`
					`}`

Alert Objects

Lockbox configuration is required for viewing alerts using the REST API. For more information about Lockbox configuration, see Configure SNMP properties after deployment on page 544, Step 1.

Table 106 Alert Object

Properties	Type	Required/ Optional
id	String	-

Table 106 Alert Object (continued)

Properties	Type	Required/ Optional
name	String	-
severity	ALERT, LOW, ALERT, MEDIUM or ALERT HIGH	-
alertType	ALERT_TYPE_UNDEFINED, DEVICE_FAILED, SDC_DISCONNECTED, MDM_NOT_CLUSTERED, SDS_DISCONNECTED, STORAGE_POOL_HAS_FAILED_CAPACITY, STORAGE_POOL_HAS_DEGRADED_CAPACITY, STORAGE_POOL_UNBALANCED, LICENSE_EXPIRED, LICENSE_ABOUT_TO_EXPIRE, FWD_REBUILD_STUCK, BKWD_REBUILD_STUCK, REBALANCE_STUCK, SDS_DISCONNECTS_FREQUENTLY, MDM_FAILS_OVER_FREQUENTLY, DRL_MODE_NON_VOLATILE, CAPACITY_UTILIZATION_ABOVE_CRITICAL_THRESHOLD, CAPACITY_UTILIZATION_ABOVE_HIGH_THRESHOLD, CONFIGURED_SPARE_CAPACITY_SMALLER_THAN_LARGEST_FAULT_UNIT, PD_INACTIVE, DEVICE_PENDING_ACTIVATION, SDS_RMCACHE_MEMORY_ALLOCATION_FAILED, NOT_ENOUGH_FAULT_UNITS_IN_SP,SDC_MAX_COUNT, FAILURE_RECOVERY_CAPACITY_BELOW_THRESHOLD, FIXED_READ_ERROR_COUNT_ABOVE_WARNING_THRESHOLD, DI_COMPARE_ERROR, PHYSICAL_DRIVE_BAD_STATE,PHYSICAL_DRIVE_HOT_SPARE, PHYSICAL_DRIVE_NOT_IN_USE, PHYSICAL_DRIVE_ENDURANCE_USED_ABOVE_THRESHOLD, BOOT_DRIVE_INVALID_STATE, LOGICAL_DISK_INVALID_READ_AHEAD_POLICY, LOGICAL_DISK_INVALID_WRITE_BACK_POLICY, LOGICAL_DISK_INVALID_ACCESS_MODE, LOGICAL_DISK_INVALID_RAID_LEVEL, LOGICAL_DISK_INVALID_CACHE_POLICY, STORAGE_CONTROLLER_INVALID_STATE, STORAGE_CONTROLLER_INVALID_TEMPERATURE, CPU_SOCKET_INVALID_VR_TEMPERATURE, CPU_SOCKET_INVALID_VR_VOLTAGE, RAM_INVALID_TEMPERATURE, RAM_INVALID_VR_TEMPERATURE, RAM_INVALID_VR_VOLTAGE, PHYSICAL_DRIVE_INVALID_TEMPERATURE, FAN_INVALID_SPEED, PSU_INVALID_INPUT, ESRS_SYSLOG, ESRS_NOT_REGISTERED, ESRS__REACHED_CAPACITY_LIMIT, NODE_INVALID_TEMPERATURE, NODE_INVALID_VOLTAGE, NODE_FAILED_TO_CONNECT_TO_BMC, SNMP_DUMMY_TRAP, STORAGE_CONTROLLER_CACHECADE_NOT_LICENSED, NODE_WITH_NO_SDC, NODE_FAILED_DEPLOYMENT,	-

Table 106 Alert Object (continued)

Properties	Type	Required/ Optional
	FIXED_READ_ERROR_COUNT_ABOVE_CRITICAL_THRESHOLD, TRIAL_LICENSE_USED, MDM_CONNECTION_LOST, NODE_FAILED_TO_CONNECT_TO_HOST, NODE_FAILED_TO_CONNECT_TO_ESX, PSU_POWER_LOST, NODE_SERIAL_NUMBER_CHANGE, PHYSICAL_DRIVE_INVALID_PD_CACHE_POLICY, XTREMCACHE_INVALID_STATE, NODE_FAILED_TO_CONNECT_TO_VCENTER, LOGICAL_DISK_NO_LONGER_CACHED, STORAGE_CONTROLLER_NOT_ALL_SLOTS_FULL, CLUSTER_DEGRADED, UPGRADE_IN_PROGRESS, CERTIFICATE_ABOUT_TO_EXPIRE, UNTRUSTED_CERTIFICATE, OBJECT_HAS_OSCILLATING_FAILURES, OBJECT_HAS_OSCILLATING_NETWORK_FAILURES, SDS_IN_MAINTENANCE, RFCACHE_DEVICE_DOES_NOT_EXIST, RFCACHE_API_ERROR_MISMACH, RFCACHE_INCONSISTENT_CACHE_CONFIGURATION, RFCACHE_INCONSISTENT_SOURCE_CONFIGURATION, RFCACHE_INVALID_DRIVER_PATH, RFCACHE_LOW_RESOURCES, RFCACHE_IO_STACK_ERROR, RFCACHE_CACHE_SKIP_DUE_TO_HEAVY_LOAD, RFCACHE_CARD_IO_ERROR, GW_CONFIGURATION_INVALID_MDM_CREDENTIALS, MDM_CERTIFICATE_EXPIRED, MDM_SECURE_CONNECTION_NOT_SUPPORTED, GW_TOO_OLD, MDM_SECURE_CONNECTION_DISABLED, MDM_CA_SIGNED_CERTIFICATE_CA_NOT_TRUSTED, MDM_SELF_SIGNED_CERTIFICATE_NOT_TRUSTED, GW_USER_REQUIRES_PW_CHANGE or MDM_CERTIFICATE_NOT_YET_VALID, MDM_CREDENTIALS_ARE_NOT_CONFIGURED, MDM_NOT_CLUSTERED_VOLUMES_EXIST, DEVICE_ERROR_ERROR, DEVICE_ERROR_WARNING, DEVICE_ERROR_NOTICE, DEVICE_ERROR_INFO, SMART_TEMPERATURE_STATE_FAILED_NOW, SMART_END_OF_LIFE_STATE_FAILED_NOW, SMART_AGGREGATED_STATE_FAILED_NOW	
startTime	Date in ISO 8601, UTC	-
lastObserved	Date in ISO 8601, UTC	-
uuid	String	-
affectedObject	Set of String describing the affected object - class name, object id...	-
alertValues		-

Table 107 Alert Operations

Operation	URI	Method	Parameters	Notes	Return
	/api/types/Alert/instances	GET			List of all Alert objects, in Alert Objects on page 512.
	/api/instances/Alert::{id}	GET			See Alert Objects on page 512

ESRS Objects

Table 108 ESRS Operations

Operation	URI	Method	Parameters	Notes	Return
Register on the ESRS server	/api/types/Esrs/actions/register	POST	Required: gateWayIp - esrs server ip siteUser - user name sitePassword - password Optional: connectInIp - special IP address that the user may give that allows remote connection to the customer site (IP address that is not blocked in FW for SSH or RDP)		
Unregister from ESRS server	/api/types/Esrs/actions/unregister	POST			
Unregister from ESRS server and remove certificate	/api/types/Esrs/actions/remove	POST	Required: gateWayIp		
Check connection	/api/types/Esrs/actions/	POST			"message" : "Esrs connected"

Table 108 ESRS Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
to an ESRS gateway	checkConnection				
Add a connect in IP address to a registered system	/api/types/Esrs/actions/setConnectInIp	POST	Required: connectInIp	This can be executed on top of register command, if it was done without connectInIp	
Start passing alerts to ESRS gateway	/api/types/Esrs/actions/start	POST			"message": "Esrs started"
Stop passing alerts to ESRS gateway	/api/types/Esrs/actions/stop	POST			"message": "Esrs stopped"

Responses

Output Content Format

The supported content format is JSON:

```
Content-Type: application/json
```

Return Codes

See http://en.wikipedia.org/wiki/List_of_HTTP_status_codes for standard HTTP status codes. In this interface, the following codes are used:

Table 109 Return codes

Status Code	Name	ScaleIO usage
200	OK	Successful GET. Body contains requested resource.
201	Created	POST has created the specified resource. Body contains minimal representation of that resource - ID and self link. (Future: etag will be in header.)
202	Accepted	Not in use yet. Will be used to return async task resource for long-running jobs.
204	No Content	Successful action POST, no body content. (Future: For modify, new etag will be in header.)

Table 109 Return codes (continued)

Status Code	Name	ScaleIO usage
400	Bad Request	Badly formed URI, parameters, headers, or body content. Essentially a request syntax error.
403	Forbidden	Not allowed - ScaleIO Gateway is disabled. Enable the gateway by editing the file `<gateway installation directory>/webapps/ROOT/WEB-INF/classes/gatewayUser.properties`. The parameter `features.enable_gateway` must be set to `true`, and then you must restart the scaleio-gateway service.
404	Not Found	Resource doesn't exist - either an invalid type name for instances list (GET, POST) or an invalid ID for a specific instance (GET, POST /action)
405	Method Not Allowed	This code will be returned if you try to use a method that is not documented as a supported method.
406	Not Acceptable	Accept headers do not meet requirements (for example, output format, version, language).
409	Conflict	The request could not be completed due to a conflict with the current state of the resource. This code is only allowed in situations where it is expected that the user might be able to resolve the conflict and resubmit the request. The response body SHOULD include enough information for the user to correct the issue.
422	Unprocessable Entity	Semantically invalid content on a POST, which could be a range error, inconsistent properties, or something similar
500	Internal Server Error	This code is returned for internal errors - file an AR. It also is returned in some platform management cases when PAPI cannot return a decent error. Best practice is to avoid filing an AR.
501	Not Implemented	Not currently used
503	Service Unavailable	This code should be returned under 2 conditions: • Apache is up but Tomcat is not, or • Apache and Tomcat are up, but Spring is not up Wait, and then try again.

For all 4xx and 5xx return codes, the body may contain an Error instance with more specifics about the failure.

Changes that might require adjustments of old code

This section describes changes that were introduced in version 2.0 that may require adjustments of code.

Deleted scsiInitiator commands

`scsiInitiator` commands were deleted. The feature is no longer supported, and was removed.

Commands:

`removeScsiInitiator`

`addMappedScsiInitiator`

`removeMappedScsiInitiator`

`ScsiInitiator`

Other deleted commands

`resetAdminPassword`

`setRemoteReadOnlyLimitState`

Cluster commands

Commands that were used to create and remove MDM clusters in versions prior to version 2.0 are no longer relevant.

Commands:

`addPrimaryMdm`

`addSecondaryMdm`

`addTiebreaker`

`removeSecondaryMdm`

`removeTiebreaker`

Instead, the following commands are now supported:

`createMdmCluster`

`addStandbyMdm`

`removeStandbyMdm`

`replaceClusterMdm`

Modified command names

The following command names changed, but when working with version 1.32 clusters, the old commands should be used.

Old command	New command
switchMdmOwnership	changeMdmOwnership
setMdmMode	switchClusterMode

Changes in command usage

`addMappedSdc –"allSdcs"` was removed.

`setManagementIps —"id"` (mdm id) was added (required field when working with a version 2.0 or higher cluster)

Changes in response format

- The "System" object format changed. When working with version 1.32 systems, "System" returns output in the old format.
- DI_COMPARE_ERROR changed to BACKGROUND_SCANNER_COMPARE_ERROR
- NODE_FAILED_TO_CONNECT changed to NODE_FAILED_TO_CONNECT_TO_BMC

Alerts name change

- Alerts exist in REST from version 1.33. The new names are also relevant also when working with ScaleIO version 1.33 systems.

Deprecated alerts

The following alerts are no longer being used as of v2.0.1:

- MDM_NOT_CLUSTERED_MANUALLY
- FAILED_CONNECT
- BACKGROUND_SCANNER_COMPARE_ERROR
- CACHECADE_USAGE_IN_STORAGE_POOL_NOT_UNIFORM
- DEVICE_SHOULD_USE_AS_CACHECADE_BUT_IT_IS_NOT
- NODE_INVALID_AIRFLOW
- NODE_SVM_OS_DISK_USED_SPACE_HIGH
- XTREM_CACHE_USAGE_IN_STORAGE_POOL_NOT_UNIFORM
- UNSTABLE_DATA_NETWORK
- STORAGE_POOL_EXTREMELY_UNBALANCED

Using the IM REST API

You can use the IM REST API to deploy a ScaleIO system. The IM is an orchestration engine that runs commands from different queues (nodes and MDM) at different "phases", and labels each one as a "process".

The IM REST API URIs begin with the “/im/” prefix.

This section includes the following topics:

General information

Most IM operations require a topology/configuration object as a payload for the REST (HTTP) request in JSON format. The following example shows a typical object in JSON format:

```
{
"snmpIp": null,
"safeIPsForLia": null,
"installationId": null,
"mdmIPs": ["10.76.60.48",
```

```
"10.76.60.41"],
"mdmPassword": "Password1",
"liaPassword": "Password1",
"licenseKey": null,
"licenseType": "unlimitedCapacity",
"isClusterOptimized": null,
"callHomeConfiguration": null,
"remoteSyslogConfiguration": null,
"systemVersionName": "",
"upgradeRunning": false,
"clustered": false,
"masterMdm": {
"node": {
"ostype": "linux",
"nodeName": null,
"nodeIPs": ["10.76.60.41"],
"domain": null,
"userName": "root",
"password": "Password1",
"liaPassword": null
},
"nodeInfo": null,
"rollbackVersion": null,
"mdmIPs": ["10.76.60.41"],
"name": "M2",
"id": null,
"ipForActor": null,
"managementIPs": ["10.76.60.41"]
},
"slaveMdmSet": [{
"node": {
"ostype": "linux",
"nodeName": null,
"nodeIPs": ["10.76.60.48"],
"domain": null,
"userName": "root",
"password": "Password1",
"liaPassword": null
},
"nodeInfo": null,
"rollbackVersion": null,
"mdmIPs": ["10.76.60.48"],
"name": "M1",
"id": null,
"ipForActor": null,
"managementIPs": ["10.76.60.48"]
}],
"tbSet": [{
"node": {
"ostype": "linux",
"nodeName": null,
"nodeIPs": ["10.76.60.146"],
"domain": null,
"userName": "root",
"password": "Password1",
"liaPassword": null
},
"nodeInfo": null,
"rollbackVersion": null,
"mdmIPs": ["10.76.60.146"],
"name": "M3",
"id": null,
"tbIPs": ["10.76.60.146"]
}],
"standbyMdmSet": [],
"standbyTbSet": [],
"sdsList": [{
"node": {
```

```
"ostype": "linux",
"nodeName": null,
"nodeIPs": ["10.76.60.146"],
"domain": null,
"userName": "root",
"password": "Password1",
"liaPassword": null
},
"nodeInfo": null,
"rollbackVersion": null,
"sdsName": "SDS_10.76.60.146",
"protectionDomain": "domain1",
"protectionDomainId": null,
"faultSet": null,
"allIPs": ["10.76.60.146"],
"sdsOnlyIPs": null,
"sdcOnlyIPs": null,
"devices": [],
"rfCached": false,
"rfCachedPools": [],
"rfCachedDevices": [],
"optimized": false,
"port": 7072,
"id": "0"
},
{
"node": {
"ostype": "linux",
"nodeName": null,
"nodeIPs": ["10.76.60.48"],
"domain": null,
"userName": "root",
"password": "Password1",
"liaPassword": null
},
"nodeInfo": null,
"rollbackVersion": null,
"sdsName": "SDS_10.76.60.48",
"protectionDomain": "domain1",
"protectionDomainId": null,
"faultSet": null,
"allIPs": ["10.76.60.48"],
"sdsOnlyIPs": null,
"sdcOnlyIPs": null,
"devices": [],
"rfCached": false,
"rfCachedPools": [],
"rfCachedDevices": [],
"optimized": false,
"port": 7072,
"id": "0"
},
{
"node": {
"ostype": "linux",
"nodeName": null,
"nodeIPs": ["10.76.60.41"],
"domain": null,
"userName": "root",
"password": "Password1",
"liaPassword": null
},
"nodeInfo": null,
"rollbackVersion": null,
"sdsName": "SDS_10.76.60.41",
"protectionDomain": "domain1",
"protectionDomainId": null,
"faultSet": null,
```

```
"allIPs": ["10.76.60.41"],
"sdsOnlyIPs": null,
"sdcOnlyIPs": null,
"devices": [],
"rfCached": false,
"rfCachedPools": [],
"rfCachedDevices": [],
"optimized": false,
"port": 7072,
"id": "0"
}],
"sdcList": [{
"node": {
"ostype": "linux",
"nodeName": null,
"nodeIPs": ["10.76.60.41"],
"domain": null,
"userName": "root",
"password": "Password1",
"liaPassword": null
},
"nodeInfo": null,
"rollbackVersion": null,
"splitterRpaIp": null,
"optimized": false
},
{
"node": {
"ostype": "linux",
"nodeName": null,
"nodeIPs": ["10.76.60.48"],
"domain": null,
"userName": "root",
"password": "Password1",
"liaPassword": null
},
"nodeInfo": null,
"rollbackVersion": null,
"splitterRpaIp": null,
"optimized": false
}],
"securityConfiguration": {
"allowNonSecureCommunicationWithMdm": false,
"allowNonSecureCommunicationWithLia": false,
"disableNonMgmtComponentsAuth": false
},
}
```


The Node object is an attribute of each ScaleIO component (for example, on which node the MDM is installed) that describes how the node will be accessed:

- via SSH (Linux node with root credentials)
- via WMI (Windows node with admin domain and credentials)
- via LIA (unknown node with null domain, null credentials and lia password)

The following samples illustrate access methods:

Linux node accessed via SSH


```
"node": {
"ostype": "unknown",
"nodeName": null,
"nodeIPs": ["10.76.60.48"],
"domain": null,
```

```
"userName": null,
"password": null,
"liaPassword": null
},
```

Unknown node (Linux or Windows) - accessed via LIA

```
"node": {
"ostype": "linux",
"nodeName": null,
"nodeIPs": ["10.76.60.41"],
"domain": null,
"userName": "root",
"password": "Password1",
"liaPassword": null
},
```

How to use IM REST API URIs

JSESSION ID must be used to in order to run IM REST API URIs. URIs run without JSESSION ID will fail.

The following steps can be used to obtain a JSESSION ID.

Procedure

1. Run the following URI:

   ```
   https://localhost/j_spring_security_check
   Content: "j_username=admin&j_password=Password1&submit=Login"
   POST
   ```

   where the user name in the above example is "admin" and the password is "Password1".

2. Find the JSESSION ID in the URI output. For example:

   ```
   JSESSIONID=969F624A761937AE80E6CC9E91756B10
   ```

3. Set the following headers, using the JSESSION ID obtained in the previous steps. For example:

   ```
   "Content-Type":"application/json"
   "Cookie":"JSESSIONID=969F624A761937AE80E6CC9E91756B10"
   ```

Results

The JSESSION ID will be added to the URIs. You may now run any IM REST API URI.

Figure 53 Example of IM REST API URI when JSESSION ID is configured as a header

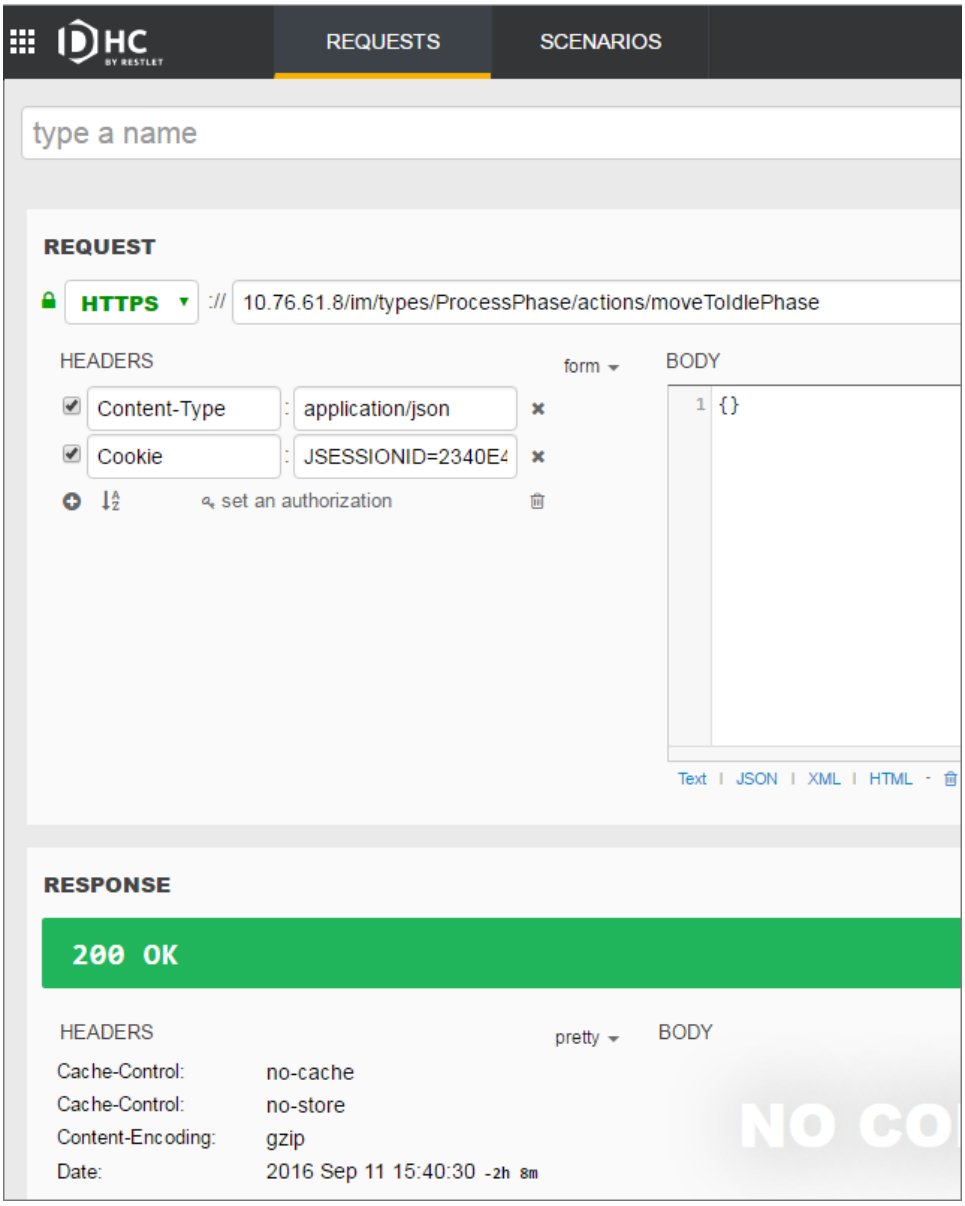

Deploying a system

When using the IM REST API to deploy a system, use JSON objects instead of the CSV file, and the appropriate steps of IM REST API commands to perform these activities:

Procedure

1. Upload the software packages
2. Install the software packages
3. Configure the system
4. Monitor progress

   Any post installation operations MUST be performed via LIA. Operations such as upgrades may fail unless all nodes have been targeted to be accessed via LIA.

To deploy the system, use the following work flow. A detailed example of a JSON configuration is shown in "General information".

5. Upload packages to the Installation Manager:

```
POST /im/types/InstallationPackages/instances/actions/
uploadPackages
```

6. Install the packages (query phase starts automatically), using the JSON configuration:

```
/im/types/Configuration/actions/install
```

7. Check progress until the installation phase is finished:

```
/im/types/Command/instances
```

8. When the installation phase is completed, move to the next phase:

```
/im/types/ProcessPhase/actions/moveToNextPhase
```

Installation Manager commands

Table 110 Operations

Operation	URI	Method	Parameters	Notes	Return

Get All IM commands in the queue	/im/types/Command/instances	GET			List of commands for current process and their status.
					Typical command:
					{
					"node": {
					/* several fields */
					},
					/* more fields */
					"commandState": "pending",
					"startTime": null,
					"message": null,
					"result": null,

Table 110 Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
					"followingCommand": "None", "logFilename": null, "archived": false, "allowedPhase": "install", "commandName": ".InstallCommand" }
Clear (flush) the IM command queues for the current process	/im/types/Command/instances/actions/clear	POST			
Abort an IM process or specific command or specific phase	/im/types/Command/instances/actions/abort	POST	At the parameter {phase: "query"} for specific phase (empty payload to abort the whole process) At the URI - add taskID for the specific command to be aborted		
Retry current phase or specific command	/im/types/Command/instances/actions/retry	POST	taskID URI long parameter to specify specific command to retry If not specified, all commands in current phase that failed / aborted will be retried		
Get current operation	/im/types/ProcessPhase/instances/	GET			{"phase":"query","nextPhase":null,"operation":"coll

Table 110 Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
and phase details					ectLogs","upgradePersistenceRecord":null,"rollbackEnabled":false } Describes the current process, phase, upgrade persistency status, and whether rollback should be allowed.
Move to next phase in process	/im/types/ProcessPhase/actions/moveToNextPhase	POST			
Move to stub phase which will mark process end (but not cleaning of the queues)	/im/types/ProcessPhase/actions/moveToIdlePhase	POST			
Begin installation or extend process	/im/types/Configuration/actions/install	POST	Parameter of Configuration URI parameters: True/False noUpload, True/False noInstall, True/False noConfigure, True/False noLinuxDevValidation, True/False extend True/False globalZeroPadPolicy		

Table 110 Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
Rollback from stuck upgrade process	/im/types/Configuration/actions/rollbackUpgrade	POST			
Begin upgrade process	/im/types/Configuration/actions/upgrade	POST	Configuration as parameter, and URI parameters: True/False noUpload, True/False noInstall, True/False noLia, True/False noWait, True/False parallelSds, Integer parallelLevel, True/False allowReboot, True/False liaOnly,		
Begin uninstall process	/im/types/Configuration/actions/uninstall	POST	Configuration as payload URI parameters: True/False noLia, True/False allowReboot,		
Set login banner to an existing cluster	/im/types/Configuration/actions/setLoginBanner	POST	Required: loginBanner - string of the login banner content mdmIPs - array of the mdm IPs	Example: {"loginBanner":"login banner string:/","mdmIPs": ["10.136.220.	

Table 110 Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			mdmPassword - string of the mdm password securityConfiguration - object contains 2 booleans - allowNonSecureCommunicationWithMdm & allowNonSecureCommunicationWithLia. Optional: None	249","10.136.220.248"] ,"mdmPassword":"Password1","securityConfiguration": {"allowNonSecureCommunicationWithMdm":false,"allowNonSecureCommunicationWithLia":false}}	
Get list of packages in the IM repository	/im/types/InstallationPackages/instances	GET	URI parameter: True/False onlyLatest		
Upload package(s) to IM repository	/im/types/InstallationPackages/instances/actions/uploadPackages	POST	Files to be uploaded - as payload in standard base64 file format		
Delete a file from IM package repository	/im/types/InstallationPackages/instances/actions/delete::{fileName}	DELETE			True/False for success
Retrieve topology from existing MDM	/im/types/Configuration/instances	POST	Parameter of: {"mdmIps": ["10.76.60.41"], "mdmUser":"admin" ,"mdmPassword":"Password1", "securityConfiguration": {"allowNonSecureCommunicationWithMdm":true,"		Configuration / topology

Table 110 Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
			allowNonSecure Communication WithLia":true,"di sableNonMgmtC omponentsAuth": false}}		
Start get info process	/im/types/ NodeInfo/ instances/ actions/ getInfo	POST	Parameter of configuration URI parameters: True/False copyRepositories , True/False collectException Only, , True/False liteVersion,		
Download zip of all previous get info results	/im/types/ NodeInfo/ instancing	GET			ZIP file
Set SSH key file	/im/types/ Configuration /actions/ setSshKeyFile	POST	Parameter of: {sshKeyFile:keyS tring} OR Payload of : {sshKeyFile:keyS tring, passphrase=phra se}	String of SSH key file is the actual key to be written to the SSH file that will be used via installation. Passphrase is the pass phrase to unlock the file	
Remove (clean) ssh key file	/im/types/ Configuration /actions/ removeSshKe yFile	POST		Disables the ability to use key file	
Get all unapproved	/types/ Upgrade/	GET			Array of certificates

Table 110 Operations (continued)

Operation	URI	Method	Parameters	Notes	Return
certificates that caused failure	instances/actions/getFailedCertificates				
Approve certificate to be based on thumbprint from the certificate which were not approved and caused failure	/types/Configuration/instances/actions/addCertificates	POST	Parameter of Thumbprint	Only thumbprints of certificates which were not approved will be added to the system. There is no way to push thumbprints or certificates that do not exist in the unapproved list.	

CHAPTER 21

SNMP trap support

This chapter describes ScaleIO support for SNMP. Topics include:

- General 534
- Supported alerts and event numbering conventions 534
- Configure SNMP properties after deployment 544
- ScaleIO.mib file 546

General

SNMP traps are implemented as part of the ScaleIO Gateway, using SNMP v2. UDP transport is used for SNMP, and the default port for trap communication is 162. The SNMP feature is disabled, by default. If you want to use the SNMP feature, enable it by editing the `gatewayUser.properties` file. For more information, see Configure SNMP properties after deployment on page 544 .

The SNMP trap sender includes a proprietary/custom `MIB` called `scaleio.mib`. This `MIB` file is located on the ScaleIO gateway server, in the `/gateway/webapps/ROOT/WEB-INF/classes` folder. A copy of the `MIB` file is included in at the end of this section. A general trap type with a unique identification number (OID) is defined in the `MIB`, so that the SNMP traps are configured to contain alert data within themselves, and use a single OID (as opposed to granular traps). All the SNMP traps contain variable bindings for severity type, which is the alert classification; the ID of the source object for which the alert was created; and an action code, which is the event number.

The alerts are calculated based on MDM polling. A repeating trap will be sent the first time that it occurs, and will only be sent again if the resend interval has passed since it was last sent. The resend frequency parameter can be configured using the **Settings** window in the ScaleIO GUI.

Only TRAP commands/messages are supported, and are initiated by the ScaleIO SNMP traps manager. GET/SET operations are not supported (or more specifically, GET/GET NEXT/GET BULK/SET/INFORM/RESPONSE).

In addition to SNMP traps, alert messages are also displayed in the GUI.

Both the ScaleIO gateway and the SNMP trap receivers must be configured. Traps can be sent to up to two SNMP trap receivers. The scaleio-gateway service must be restarted after configuration.

Supported alerts and event numbering conventions

The following alerts can be sent as SNMP traps by the ScaleIO system. All events are numbered in the following format: SIO<CLASS>.<TYPE>.<ISSUE>. The issue number is a running counter for all issues in a specific type.

Open and closing alerts will consist of the same code and issue number, with the exception of the first digit (0 or 1) in the <ISSUE> section. For example:

- SIOXX.XX.0XXXXXX indicates that the alert is active
- SIOXX.XX.1XXXXXX indicates that the alert has been closed

CLASS/TYPE:

System = 1

- Capacity = 1
- License = 2

MDM = 2

- MDM_Cluster = 1
- Protection_Domain = 2
- Fault_Set = 3

- Storage_Pool = 4

SDS = 3

- SDS = 1
- Device = 2

SDC = 4

- SDC = 1

Volume = 5

ESRS = 10

- ESRS = 1

Note

Each alert has a corresponding Closed state, represented by the code SIOXX.XX.1XXXXXX. For example Open state: SIO02.01.0000003, Closed state: SIO02.01.1000003

ScaleIO Alerts in SNMP, GUI, REST, and ESRS

This table summarizes alerts generated by ScaleIO systems.

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
License expired	LICENSE_EXPIRED	System.License.License_Expired	SIO01.02.0000001	5 (Critical)	To resume operational mode, contact EMC Support for license renewal. If you have already renewed your license, install it.
The system's license will expire in n days	LICENSE_ABOUT_TO_EXPIRE	System.License.License_Is_About_To_Expire	SIO01.02.0000002	3 (Error) 2 (Warning) according to time left and limits	Contact EMC Support for license renewal. If you have already renewed your license, install it.
ScaleIO is using a trial license	TRIAL_LICENSE_USED	System.License.Trial_License_Used	SIO01.02.0000003	2 (Warning)	Purchase a license and install it.
Oscillating failures reported	OBJECT_HAS_OSCILLATING_FAILURES	System.Oscillating_Failures.Object_has_oscillating_failures	SIO01.03.0000001	2 (Warning)	Check oscillating failures of the component and take action accordingly. If the oscillating failure does not indicate a problem, change the settings of the oscillating failure window to supress this alert
There are oscillating network failures	OBJECT_HAS_OSCILLATING_NETWORK_FAILURES	System.Oscillating_Failures.OBJECT_HAS_OSCILLATING_NETWORK_FAILURES	SIO01.03.0000002	2 (Warning)	Check the oscillating failure report, that can be accessed from one of the management interfaces. Check whether there is a

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS (continued)

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
					problem with network links, fix, and restart the counters.
No valid MDM credentials are configured in ScaleIO Gateway	GW_CONFIGURATION_INVALID_MDM_CREDENTIALS	System.Credentials.GW_CONFIGURATION_INVALID_MDM_CREDENTIALS	SIO01.04.0000001	5 (Critical)	Configure the MDM credentials in the ScaleIO Gateway using the SioGWTool.
MDM credentials are not configured in the ScaleIO Gateway	MDM_CREDENTIALS_ARE_NOT_CONFIGURED	System.Credentials.MDM_CREDENTIALS_ARE_NOT_CONFIGURED	SIO01.04.0000002	5 (Critical)	Configure MDM credentials on the ScaleIO Gateway using the SioGWTool
The MDM user configured in ScaleIO Gateway requires a password change	GW_USER_REQUIRES_PW_CHANGE	System.Credentials.GW_USER_REQUIRES_PW_CHANGE	SIO01.04.0000004	5 (Critical)	Configure MDM credentials on the ScaleIO Gateway using the SioGWTool
System upgrade is in progress	UPGRADE_IN_PROGRESS	System.Upgrade.UPGRADE_IN_PROGRESS	SIO01.05.0000001	3 (Error)	Monitor the upgrade process, and check that it is completed successfully.
ScaleIO Gateway version is too old	GW_TOO_OLD	System.Upgrade.GW_TOO_OLD	SIO01.05.0000002	5 (Critical)	Upgrade the ScaleIO Gateway to the same version as the rest of your system.
The MDM is not operating in clustered mode	MDM_NOT_CLUSTERED	MDM.MDM_Cluster.MDM_Not_Clustered	SIO02.01.0000001	5 (Critical)	MDM cluster was manually set to SINGLE mode. Confirm that this is an expected operation. Working in SINGLE mode is not recommended. Prepare the cluster modules (if needed), and return to CLUSTER mode.
MDM fails over frequently	MDM_FAILS_OVER_FREQUENTLY	MDM.MDM_Cluster.MDM_Fails_Over_Frequently	SIO02.01.0000003	5 (Critical) 3 (Error) 2 (Warning) according to disconnect count and hardcoded values (2/3/10)	The MDMs frequently swap ownership. No action required.
Forward rebuild cannot proceed	FWD_REBUILD_STUCK	MDM.MDM_Cluster.FWD_REBUILD_STUCK	SIO02.01.0000004	2 (Warning)-5 (Critical)	Check the system for lack of spare capacity and/or failed capacity, and either fix the problem or add capacity if necessary.

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS (continued)

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Backward rebuild cannot proceed	BKWD_REBUILD_STUCK	MDM.MDM_Cluster.BKWD_REBUILD_STUCK	SIO02.01.0000005	2 (Warning)-5 (Critical)	Check the system for lack of spare capacity and/or failed capacity, and either fix the problem or add capacity if necessary.
Rebalance cannot proceed	REBALANCE_STUCK	MDM.MDM_Cluster.REBALANCE_STUCK	SIO02.01.0000006	5 (Critical) 3 (Error) 2 (Warning)	Add a physical disk; if this is not possible, reduce the spare policy while maintaining enough spare to sustain a rebuild, if necessary.
The MDM cluster is degraded, and data is not protected	CLUSTER_DEGRADED	MDM.MDM_Cluster.CLUSTER_DEGRADED	SIO02.01.0000007	3 (Error)-5 (Critical)	Check that all MDM cluster nodes are functioning correctly, and fix and replace faulty nodes, if necessary, in order to return to full protection.
	MDM_CONNECTION_LOST	MDM.MDM_Cluster.MDM_CONNECTION_LOST	SIO02.01.0000008		Check the connection to the MDM
The MDM is not operating in Clustered mode	MDM_NOT_CLUSTERED_VOLUMES_EXIST	MDM.MDM_Cluster.MDM_Not_Clustered_Volume_Exist	SIO02.01.0000009	5 (Critical)	The MDM cluster was manually set to SINGLE mode. Working in SINGLE mode is not recommended. Single mode means that there is only one copy of the MDM repository. If you lose this copy, all System configurations and all the data on all the existing volumes will be lost. Please verify that this is an expected operation. Prepare the cluster modules (if needed), and return to CLUSTERED mode as soon as possible.
Inactive Protection Domain	PD_INACTIVE	MDM.Protection_Domain.Protection_Inactive	SIO02.02.0000001	2 (Warning)	Protection Domain was inactivated by a user command. Confirm that this is an expected operation. This is usually done for maintenance. When maintenance is complete, reactivate the Protection Domain.

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS (continued)

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Storage Pool has failed capacity	STORAGE_POOL_HAS_FAILED_CAPACITY	MDM.Storage_Pool.Storage_Pool_has_Failed_Capacity	SIO02.04.0000001	5 (Critical)	For the given Storage Pool, for some blocks, both primary and secondary copies are inaccessible. Check and fix the state of all devices in the Storage Pool and all the server's holding devices in the Storage Pool.
Storage Pool has degraded capacity	STORAGE_POOL_HAS_DEGRADED_CAPACITY	MDM.Storage_Pool.Storage_Pool_has_Degraded_Capacity	SIO02.04.0000002	3 (Error)	For the given Storage Pool, for some blocks, one of the two copies is inaccessble. Check if a server is offline or if there is another server hardware-related issue. Check if a storage device is down.
Capacity utilization above critical threshold	CAPACITY_UTILIZATION_ABOVE_CRITICAL_THRESHOLD	MDM.Storage_Pool.Capacity_Utilization_Above_Critical_Threshold	SIO02.04.0000003	5 (Critical)	Due to thinly provisioned volumes or snapshot usage, the capacity utilization of the Storage Pool is reaching a critical threshold. Remove snapshots, if possible, or add physical storage.
Capacity utilization above high threshold	CAPACITY_UTILIZATION_ABOVE_HIGH_THRESHOLD	MDM.Storage_Pool.Capacity_Utilization_Above_High_Threshold	SIO02.04.0000004	3 (Error)	Due to thinly provisioned volumes or snapshot usage, the capacity utilization of the Storage Pool is reaching a high threshold. Remove snapshots, if possible, or add physical storage.
Capacity utilization above high threshold	CAPACITY_UTILIZATION_ABOVE_HIGH_THRESHOLD	MDM.Storage_Pool.Capacity_Utilization_Above_High_Threshold	SIO02.04.0000004	2 (Warning)	Due to thinly provisioned volumes or snapshot usage, the capacity utilization of the Storage Pool is reaching a high threshold. Remove snapshots, if possible, or add physical storage.
Failure recovery capacity is below the threshold	FAILURE_RECOVERY_CAPACITY_BELOW_THRESHOLD	MDM.Storage_Pool.Failure_Recovery_Capacity_Below_Threshold	SIO02.04.0000005	3 (Error)	The capacity available for recovery in a degraded storage event is lower than the predefined threshold. Replace failed hardware or add more physical storage.
Configured spare capacity is smaller than largest fault unit	CONFIGURED_SPARE_CAPACITY_SMALLER_THAN_LARGEST_FAULT_UNIT	MDM.Storage_Pool.Configured_Spare_Capacity_Smaller_Tha	SIO02.04.0000008	2 (Warning)	Increase the "spare percentage" configured in the system for the Storage Pool, so that the capacity

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS (continued)

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
		n_Largest_Fault_Unit			reserved for failure recovery is larger than the largest fault unit in the Storage Pool.
The Storage Pool relies too heavily (over 50%) on capacity from a single SDS or Fault Set. Balance capacity over other SDSs or Fault Sets.	STORAGE_POOL_UNBALANCED	MDM.Storage_Pool.STORAGE_POOL_UNBALANCED	SIO02.04.0000009	3 (Error)	Move some physical disks from the large SDS to the others, or add disks to the smaller SDS in order to approximate the capacity of the large SDS as much as possible.
Storage Pool does not meet the minimum requirement of 3 fault units	NOT_ENOUGH_FAULT_UNITS_IN_SP	MDM.Storage_Pool.Not_Enough_Fault_Units	SIO02.04.0000010	3 (Error)	Add more SDSs to the Storage Pool to meet the minimum requirement of 3 hosts.
There are cluster certificates pending approval. For more information, open System Settings > Certificates.	UNTRUSTED_CERTIFICATE	MDM.CERTIFICATE.UNTRUSTED_CERTIFICATE	SIO02.05.0000001	3 (Error)	Check and the certificates and trust them you see fit. For more information, open System Settings > Certificates.
Master MDM Certificate is about to expire	CERTIFICATE_ABOUT_TO_EXPIRE	MDM.CERTIFICATE.CERTIFICATE_ABOUT_TO_EXPIRE	SIO02.05.0000002	5 (Critical)	Install a valid SSL certificate on the MDM before the old one expires
Master MDM Certificate has expired	MDM_CERTIFICATE_EXPIRED	MDM.CERTIFICATE.MDM_CERTIFICATE_EXPIRED	SIO02.05.0000003	5 (Critical)	Install a valid SSL certificate on the host
Secure connection disabled on MDM	MDM_SECURE_CONNECTION_DISABLED	MDM.CERTIFICATE.MDM_SECURE_CONNECTION_DISABLED	SIO02.05.0000004	5 (Critical)	Enable secure connections on the MDM in order to protect your login information
The self-signed certificate presented by the Master MDM is not trusted	MDM_SELF_SIGNED_CERTIFICATE_NOT_TRUSTED	MDM.CERTIFICATE.MDM_SECURE_CONNECTION_DISABLED	SIO02.05.0000005	5 (Critical)	Check the certificate, and trust it if you see fit to do so
MDM does not support secure connections	MDM_SECURE_CONNECTION_NOT_SUPPORTED	MDM.CERTIFICATE.MDM_SECURE_CONNECTION_NOT_SUPPORTED	SIO02.05.0000006	5 (Critical)	Check MDM cluster nodes
The validity period of the certificate presented by the	MDM_CERTIFICATE_NOT_YET_VALID	MDM.CERTIFICATE.MDM_CERTIFICATE_NOT_YET_VALID	SIO02.05.0000007	5 (Critical)	The time and date on the computer where the certificate was created is

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS (continued)

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Master MDM starts in the future					not consistent with the time and date set in the ScaleIO system. Replace the certificate or fix the system time.
The Certificate Authority that signed the Master MDM's certificate is not trusted	MDM_CA_SIGNED_CERTIFICATE_CA_NOT_TRUSTED	MDM.CERTIFICATE.MDM_CA_SIGNED_CERTIFICATE_CA_NOT_TRUSTED	SIO02.05.0000008	5 (Critical)	Trust the CA certificate if you see fit
SDS is disconnected	SDS_DISCONNECTED	SDS.SDS.SDS_Disconnected	SIO03.01.0000001	3 (Error)	The SDS service may be down or unreachable over the network. Verify that the SDS service is up and running and that the network is properly connected.
SDS disconnects frequently	SDS_DISCONNECTS_FREQUENTLY	SDS.SDS.SDS_Disconnected_Frequently	SIO03.01.0000002	3 (Error) 2 (Warning) according to disconnect count and hard-coded	The SDS connection is fluctuating due to an unstable network connection. Check the SDS data network connection for Packet Drops, and try to disconnect one of the ports to see if the SDS disconnection issue is resolved by using only one port. If this does not resolve the issue, switch to the other port. If there is still an issue, it may be due to a Faulty NIC, Faulty Switch ports, or a faulty switch. If there is no issue with another switch, the issue was switch-related. Otherwise, the issue may be due to a faulty NIC, which requires NIC replacement.
Memory allocation for RAM Read Cache failed on SDS	SDS_RMCACHE_MEMORY_ALLOCATION_FAILED	SDS.SDS.SDS_Rmcache_Memory_allocation_Failed	SIO03.01.0000003	2 (Warning)	The system failed to allocate memory to the SDS RAM Read-Cache. For 32 GB RAM or less, up to 50% of the memory can be allocated for caching. From 32 GB or more, up to 75% of the memory can be allocated for caching. Reduce the

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS (continued)

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
					configured RAM Read-Cache memory to match the allocation conditions.
DRL mode: Hardened	DRL_MODE_NON_VOLATILE	SDS.SDS.DRL_MODE_NON_VOLATILE	SIO03.01.0000004	1 (Info)	DRL Mode is configured to "Hardened" instead of "Volatile". Both modes are configurable.
RFcache card I/O error	RFCACHE_CARD_IO_ERROR	SDS.SDS.RFCACHE_CARD_IO_ERROR	SIO03.01.0000005	2 (Warning)	Disable caching on the device and check the health of the device, because it may be faulty. If necessary, replace the device.
RFcache skipped due to heavy load	RFCACHE_CACHE_SKIP_DUE_TO_HEAVY_LOAD	SDS.SDS.RFCACHE_CACHE_SKIP_DUE_TO_HEAVY_LOAD	SIO03.01.0000006	2 (Warning)	Read Flash Cache is working under a heavy load, and therefore has skipped some IOs. This is a temporary error which should resolve itself. If it persists, try to balance the Storage Pool contents across more SDSs, or add more cache cards.
RFcache IO stack error	RFCACHE_IO_STUCK_ERROR	SDS.SDS.RFCACHE_IO_STUCK_ERROR	SIO03.01.0000007	2 (Warning)	IO has become stuck on the cache device. Disable caching on the device and check the health of the device, because it may be faulty. If necessary, replace the device.
RFcache resources are low	RFCACHE_LOW_RESOURCES	SDS.SDS.RFCACHE_LOW_RESOURCES	SIO03.01.0000008	2 (Warning)	There is not enough RAM available on the server for Read Flash Cache optimal operation. Increase the amount of available RAM.
RFcache driver path is invalid	RFCACHE_INVALID_DRIVER_PATH	SDS.SDS.RFCACHE_INVALID_DRIVER_PATH	SIO03.01.0000009	2 (Warning)	The Read Flash cache driver (xcache) is either not installed, or was installed in the wrong location. Install the driver, and contact Customer Support if the problem persists.
RFcache source configuration is inconsistent	RFCACHE_INCONSISTENT_SOURCE_CONFIGURATION	SDS.SDS.RFCACHE_INCONSISTENT_SOURCE_CONFIGURATION	SIO03.01.0000010	2 (Warning)	Check RFcache state of all disks in the pool and adjust them so that all disks have the same caching state.

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS (continued)

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
RFcache source configuration is inconsistent	RFCACHE_INCONSISTENT_CACHE_CONFIGURATION	SDS.SDS.RFCACHE_INCONSISTENT_CACHE_CONFIGURATION	SIO03.01.0000011	2 (Warning)	Query the system to determine what is not consistent in the configurations of the Read Flash cache driver and the SDS where the cache device is located.
RFcache device does not exist	RFCACHE_DEVICE_DOES_NOT_EXIST	SDS.SDS.RFCACHE_DEVICE_DOES_NOT_EXIST	SIO03.01.0000012	2 (Warning)	You tried to add a cache device that does not exist. Check and fix Read Flash Cache configuration.
RFcache API mismatch	RFCACHE_API_ERROR_MISMACH	SDS.SDS.RFCACHE_API_ERROR_MISMACH	SIO03.01.0000013	2 (Warning)	The Read Flash Cache (xcache) driver version and SDS version do not match. Try to upgrade them to the same version. If the problem persists, contact Customer Support.
SDS is in Maintenance Mode	SDS_IN_MAINTENANCE	SDS.SDS.SDS_IN_MAINTENANCE	SIO03.01.0000014	2 (Warning)	The SDS is currently in Maintenance Mode. Exit Maintenance Mode once maintenance is complete. If an NDU is in progress, ignore this warning.
Device failed	DEVICE_FAILED	SDS.Device.Device_Failed	SIO03.02.0000001	3 (Error)	The SDS device could not be opened, read from or written to. Validate the device state. Check the cause of the error, and determine if it is a human error or a system malfunction. Check hardware if needed.
Device test is done and device is pending activation	DEVICE_PENDING_ACTIVATION	SDS.Device.Device_Pending_Activation	SIO03.02.0000002	2 (Warning)	The SDS device has been added and tested. Activate the SDS device.
Device has fixed read errors	FIXED_READ_ERROR_COUNT_ABOVE_WARNING_THRESHOLD	SDS.Device.FIXED_READ_ERROR_COUNT_ABOVE_WARNING_THRESHOLD	SIO03.02.0000003	3 (Error) if counter > 0	Read from the SDS device failed. Data was corrected from the other copy. No action is required, but note that the device might be faulty.
Device has fixed read errors	FIXED_READ_ERROR_COUNT_ABOVE_CRITICAL_THRESHOLD	SDS.Device.FIXED_READ_ERROR_COUNT_ABOVE_CRITICAL_THRESHOLD	SIO03.02.0000004	5 (Critical) if counter >= 5	SDS device read failed more than 5 times. Replace the physical device.

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS (continued)

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
Device failed: All IO to the device will be stopped, and data will be relocated to another device.	DEVICE_ERROR_ERROR	SDS.Device.Device_ErrorError	SIO03.02.0000005	5 (Critical)	Check the device, and if necessary, replace it
Device failed: All IO to the device will be stopped, and data will be relocated to another device.	DEVICE_ERROR_WARNING	SDS.Device.Device_ErrorWarning	SIO03.02.0000006	3 (Error)	Check the device, and if necessary, replace it
Device malfunction has been detected.	DEVICE_ERROR_NOTICE	SDS.Device.Device_ErrorNotice	SIO03.02.0000007	2 (Warning)	Check the device, and if necessary, replace it
Minor failures have been detected in device performance.	DEVICE_ERROR_INFO	SDS.Device.Device_ErrorInfo	SIO03.02.0000008	2 (Warning)	Check the device, and if necessary, replace it
Disk temperature is above the configured threshold, and may fail soon if no action is taken.	SMART_TEMPERATURE_STATE_FAILED_NOW	SDS.Device.SMART_Temperature_State_failed.Now	SIO03.02.0000009	3 (Error)	Check the temperature in the server and at the data center. Check if a fan alert is raised in the node.
The disk is near the end of its working life, and should be replaced.	SMART_END_OF_LIFE_STATE_FAILED_NOW	SDS.Device.SMART_End_Of_Life_State_Failed_Now	SIO03.02.0000011	3 (Error)	Replace the disk.
The disk may be about to fail, or may be operating with reduced performance.	SMART_AGGREGATED_STATE_FAILED_NOW	SDS.Device.SMART_Aggregated_State_Failed_Now	SIO03.02.0000013	3 (Error)	Consider replacing the disk.
The SDC is either down or unreachable over the network	SDC_DISCONNECTED	SDC.SDC.SDC_DISCONNECTED	SIO04.01.0000001	3 (Error)	Verify that the SDC service is up and running and that the network is properly configured and connected.
No more SDCs can be defined on this system; the maximum has been reached	SDC_MAX_COUNT	SDC.SDC.SDC_Max_Count	SIO04.01.0000002	3 (Error)	The maximum number of SDCs in the system has been reached (1024^2).
Unable to connect to ESRS Gateway	ESRS_CONNECTIVITY_ERROR	Esrs.Esrs.CONNECTIVITY_ERROR	SIO10.01.0000001	5 (Critical)	Check the network's connectivity to the ESRS Gateway.

Table 111 ScaleIO Alerts in SNMP, GUI, REST, and ESRS (continued)

Alert Message in GUI	Alert Message in REST	Alert Message in SNMP Trap	Alert Code (for ESRS)	Severity	Recommended Action
The system is not registered with an ESRS gateway	ESRS_NOT_REGISTERED	Esrs.Esrs.NOT_REGISTERED	SIO10.01.0000004	2 (Warning)	The system is not registered with an ESRS Gateway and will not send Alerts to EMC for monitoring. Contact your Support \ Sales representative to get an EMC ESRS support package.
ESRS number of messages received has been reached. No more messages will be sent in the next 8 hours	ESRS_REACHED_CAPACITY_LIMIT	Esrs.Esrs.REACHED_CAPACITY_LIMIT	SIO10.01.0000005	5 (Critical)	ESRS has a limit of receiving up to 200 alerts per 8 hours. The limit has been reached, so no ESRS messages will be sent in the following 8 hours.

Configure SNMP properties after deployment

Configure SNMP trap properties after deployment. These procedures are mandatory for VMware-based systems where the SNMP feature is required. For other operating systems, configuration can be done either during deployment, or afterwards, using the instructions in this section.

The following procedures are required to enable the SNMP feature:

1. Create a Lockbox on page 544
2. Configuring SNMP after deployment on page 545

Create a Lockbox

Create a Lockbox and add the MDM credentials. Lockbox is required for the following features: ESRS, SNMP, LDAPS.

Before you begin

The following items are required for the ESRS feature:

- ESRS Gateway v3 version 3.08 or higher must be installed and configured. It is recommended to create at least two ESRS Gateways and define them as cluster via the backend server.
- ESRS Gateway must be reachable from ScaleIO on port 9443.
- The ScaleIO license must be installed.

Ensure you have:

- One or more IP addresses of the ESRS gateway servers. Note that ESRS does not currently support IPv6.
- ESRS username and password.
- ScaleIO Gateway IP address, username, and password.
- MDM username and password.

- The ScaleIO Management IP address to be used as the Connect-In IP address. It must be an IP address that is accessible from the ESRS Gateway (for example, in case of NAT).

Use SioGWTool to configure a Lockbox. SioGWTool should be used to create a Lockbox only when a Lockbox has not yet been created.

A Lockbox can be created during installation with the Installation Manager (IM). For more information on creating a Lockbox during installation, see the *Deployment Guide*.

To use SioGWTool, input the appropriate path, based on your operating system, and append the commands to the end of the filepath:

- Linux SioGWTool filepath: `/opt/emc/scaleio/gateway/bin/SioGWTool.sh`
- Windows SioGWTool filepath: `C:\Program Files\EMC\ScaleIO\Gateway\bin\SioGWTool.bat`

Procedure

1. Create a Lockbox:

```
<SioGWTool_PATH> --change_lb_passphrase --new_passphrase
<NEW_PASSPHRASE>
```

Windows example:

```
 C:\Program Files\EMC\ScaleIO\Gateway\bin\SioGWTool.bat --
set_mdm_credentials --mdm_user admin --mdm_password Scaleio123
```

2. Add MDM credentials to the Lockbox:

```
<SioGWTool_PATH> --set_mdm_credentials --mdm_user
<MDM_USERNAME> --mdm_password <MDM_PASSWORD>
```

Configuring SNMP after deployment

Configure Simple Network Management Protocol (SNMP) for error reporting, if it was not configured during installation.

Before you begin

Ensure that a LockBox has already been created and the MDM credentials have been added to it. For detailed instructions on creating a LockBox, see Create a Lockbox on page 544.

Enable the SNMP feature in the `gatewayUser.properties` file.

Procedure

1. Use a text editor to open the `gatewayUser.properties` file, located in the following directory on the Installation Manager/Gateway server:
   - Linux: `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes`
   - Windows: `C:\Program Files\EMC\ScaleIO\Gateway\webapps\ROOT\WEB-INF\classes\`

2. Locate the parameter `features.enable_snmp` and edit it as follows:

```
features.enable_snmp=true
```

3. To add the trap receiver IP address, edit the parameter `snmp.traps_receiver_ip` as follows:

```
snmp.traps_receiver_ip <TRAP_IP_1>, <TRAP_IP_2>
```

   The SNMP trap receivers' IP address parameter supports up to two comma-separated or semi-colon-separated hostnames or IP addresses.

4. You can optionally change the following parameters:

Option	Description
snmp.sampling_frequency	The MDM sampling period. The default is 30.
snmp.resend_frequency	The frequency of resending existing traps. The default is 0, which means that traps are sent all the time.

5. Save and close the file.
6. Restart the scaleio-gateway service:
   - Linux: Run the command `service scaleio-gateway restart`
   - Windows: Restart the EMC Gateway service.

ScaleIO.mib file

The following text is the content of the `scaleio.mib` file.

Note

The object source identifier in the trap MDM.MDM_Cluster.SNMP_Server_Cannot_Connect_To_MDM is "NA".

```
SCALEIO-MIB DEFINITIONS ::= BEGIN IMPORTS MODULE-IDENTITY,
OBJECT-TYPE, NOTIFICATION-TYPE, Integer32 FROM SNMPv2-SMI
DisplayString FROM RFC1213-MIB OBJECT-GROUP, NOTIFICATION-GROUP
FROM SNMPv2-CONF emc FROM EMC-MIB;
```

```
scaleio MODULE-IDENTITY
LAST-UPDATED "201511060000Z"
ORGANIZATION "EMC Corporation"
CONTACT-INFO
"EMC Corporation

www.emc.com"

DESCRIPTION
"The Structure of Management Information for the EMC SCALEIO
enterprise."
REVISION      "201511060000Z"
```

```
DESCRIPTION
"Initial version of this MIB."
::= { emc 101 }
```

```
-- 1.3.6.1.4.1.1139.101.1
scaleioAlert OBJECT IDENTIFIER ::= { scaleio 1 }
```

```
-- 1.3.6.1.4.1.1139.101.1.1
scaleioAlertSeverity OBJECT-TYPE
SYNTAX Integer32
MAX-ACCESS accessible-for-notify
STATUS current
DESCRIPTION
"Severity of the event"
::= { scaleioAlert 1 }
```

```
-- 1.3.6.1.4.1.1139.101.1.2
scaleioAlertType OBJECT-TYPE
SYNTAX DisplayString (SIZE (0..255))
MAX-ACCESS accessible-for-notify
STATUS current
DESCRIPTION
"Type of the alert"
::= { scaleioAlert 2 }
```

```
-- 1.3.6.1.4.1.1139.101.1.3
scaleioAlertSourceObjectId OBJECT-TYPE
SYNTAX DisplayString (SIZE (0..255))
MAX-ACCESS accessible-for-notify
STATUS current
DESCRIPTION
"Object id for which the alert was created"
::= { scaleioAlert 3 }
```

```
-- 1.3.6.1.4.1.1139.101.1.4
scaleioAlertActionCode OBJECT-TYPE
SYNTAX DisplayString (SIZE (0..255))
MAX-ACCESS accessible-for-notify
STATUS current
DESCRIPTION
"Action code of the alert"
::= { scaleioAlert 4 }
```

```
-- 1.3.6.1.4.1.1139.101.1.5
scaleioGroups OBJECT IDENTIFIER ::= { scaleioAlert 5 }
```

```
-- 1.3.6.1.4.1.1139.101.1.5.1
currentObjectGroup OBJECT-GROUP
OBJECTS { scaleioAlertSeverity,
scaleioAlertType,
```

```
scaleioAlertSourceObjectId,
scaleioAlertActionCode }
STATUS current
DESCRIPTION
"scaleio-MIB-V2 OBJECT-GROUP."
::= { scaleioGroups 1 }
```

```
-- 1.3.6.1.4.1.1139.101.1.5.2
currentNotificationGroup NOTIFICATION-GROUP
NOTIFICATIONS { scaleioAEAlert }
STATUS current
DESCRIPTION
"scaleio-MIB-V2 NOTIFICATION-GROUP."
::= { scaleioGroups 2 }
```

```
scaleioAEAlert NOTIFICATION-TYPE
OBJECTS { scaleioAlertSeverity, scaleioAlertType,
scaleioAlertSourceObjectId, scaleioAlertActionCode }
STATUS current
DESCRIPTION "ScaleIO Alert"
::= { scaleio 2 }
END
```

GLOSSARY

A

Active Directory Active Directory (AD) provides directory-based identity-related services. It maintains a directory that is used to centrally store identity information and security principles, and uses them to authenticate and authorize users and devices.

Active Forward Rebuild A copy of stored data is currently being rebuilt on another server, due to planned or unplanned shutdown of a server.

B

Backward Rebuild Data is rebuilt on servers that went offline and became active again. Forward rebuilds can take a long time, and therefore, it can be quicker to restore and update the data on a server which has come back online, than it is to do an entire rebuild on a different server.

BWC Bandwidth counters.

C

Cache Cache is random access electronic storage used to retain frequently used data for faster access by the channel. Cache is a critical aspect of storage performance. ScaleIO uses server DRAM for Read RAM Cache (RMcache) as well as SSD/Flash devices (RFcache) for caching reads. ScaleIO cache uses recently-accessed (LRU) data readily available to manage caching. I/Os read from cache have a lower response time than I/Os serviced by the drives. In addition, cached I/Os reduce the data drive workload, which in many cases is a performance bottleneck in the system.

CacheCade Read and Write caching of storage devices performed by one or more designated SSD devices in the ScaleIO system.

Cache Hit Rate The percentage of I/Os from cache.

Cache Skip Data is written directly to storage, bypassing the cache. Reasons for cache skips include: I/Os were too large, the cache device was busy, or I/Os were unaligned. The cache can also be configured to always work in passthrough mode.

Cache Writes Handling Mode The caching write-mode used by the system: passthrough mode (writes to storage only), or cached mode (by default, writes both to cache and to storage).

Cluster Mode ScaleIO is controlled by a cluster of MDM nodes, minimally consisting of a Master MDM, Slave MDM, and a Tie Breaker node. 5-node clusters consist of one Master MDM, two Slave MDMs, and two Tie Breakers.

D

Degraded Capacity The capacity is available, but is not protected in case of another failure

Device Physical storage device, such as a flash drive, or magnetic disk

DirectPath In ScaleIO documentation, we use the term DirectPath to refer to the VMware vSphere VMDirectPath I/O feature.

DRL Dirty Region Logging: DRL bits indicate if data is in-writing to a certain location. Once the data is written in both primary and secondary locations, the DRL bit associated with the written location is cleared. These bits can be either stored in DRAM only (memory_only) or also backed up in non-volatile memory (hardened). The former delivers better I/O performance; the latter reduces data movement following a power-cycle giving rise to a faster rebuild.

F

Failed Capacity The capacity is inaccessible due to a failure, and data integrity is at risk

Fault Sets A logical entity that ensures that SDS data is backed up on SDSs that belong to other Fault Sets, thus preventing double-point-of-failure scenarios if rack power outages occur.

Forward Rebuild Data in storage will be rebuilt on another server, due to planned or unplanned shutdown of a server.

I

ID Identifier, a unique sequence of characters that identifies an object in the system. In some CLI commands, an ID can be used to specify a system component.

IP Role The role of the IP address configured for an SDS. Each SDS can have several IP addresses associated with it. Each IP address can serve a different purpose, or role. IP roles include: SDS, SDC, or both SDS and SDC.

L

LDAP The Lightweight Directory Access Protocol (LDAP) is a directory service protocol that runs on a layer above the TCP/IP stack. It provides a mechanism used to connect to, search, and modify Internet directories using Client-Server architecture. In ScaleIO, LDAP is the protocol used by the MDM to communicate with Active Directory (AD) for authentication purposes.

Lockbox Lockbox is a component of the RSA Common Security Toolkit (CST) which securely stores data (such as passwords) in an encrypted file. A lockbox must be defined for LDAP (secure LDAP), SNMP, and ESRS. For LDAP, lockbox use is optional.

M

Management IPs The IP addresses of the MDMs defined in the system that can be used to access the MDM from CLI, GUI and REST.

Management Port The Port number used by the MDM for purposes of communicating with the nodes in the ScaleIO network.

Manager MDM An MDM that can act as a Master or a Slave in the cluster. Manager MDMs have a unique system ID, and can be given unique names. A manager can be a standby or a member of the cluster.

Master MDM The MDM in the cluster that controls the SDSs and SDCs.

MDM Any server with the MDM package installed on it. An MDM can be given a Manager or a Tie Breaker (default) role, during installation. MDMs have a unique MDM ID, and can be given unique names.

P

Page Size The page size, typically in KB, used for caching purposes by Read Flash Cache.

Pass-Through Mode Data is passed through to or from storage devices without being cached by Read Flash Cache.

Pending Backward Rebuild A backward rebuild is waiting in a queue, and will be performed when possible, according to rebuild throttling policy.

Primary MDM See Master MDM.

Protected Capacity Capacity that has an accessible copy in the system, in case of failure.

Protection Domain A unique set of SDSs grouped together for reliability and tenancy separation.

R

RAM Read Cache (RMcache) Server RAM that is reserved for caching storage devices in a Storage Pool.

Read Flash Cache (RFcache) Read-only caching of storage devices performed by one or more designated SSD devices and PCIe flash devices in a ScaleIO system.

Rebalance When ScaleIO detects lopsided use of storage capacity, or when new nodes are added, it redistributes data across the nodes, in order to improve performance.

Rebuild When ScaleIO detects a failure in the network, it creates a new copy of the data from the failed component, in a new location, to ensure data integrity.

Restricted MDM Mode A mode set in which commands can only be performed from an MDM machine.

Restricted SDC Mode Only approved SDCs can access the MDM. When this mode is enabled, volumes can only be added to approved SDCs.

S

SDBG The ScaleIO Debugger is a ScaleIO tech support troubleshooting tool, used to investigate for "live" systems that retrieves internal information from different ScaleIO components.

SDC ScaleIO Data Client, a lightweight device driver that exposes ScaleIO volumes as block devices to the application residing on the same server on which the SDC is installed.

SDS ScaleIO Data Server, which manages the capacity of a single server and acts as a back-end for data access. The SDS is installed on all servers contributing storage devices to the ScaleIO system.

Secondary MDM See Slave MDM.

Single Mode A single MDM manages the ScaleIO network. This mode has no backup protection, and should not be used in production environments.

Slave MDM An MDM in the cluster that is ready to take over the Master MDM role if ever necessary.

Snapshot Capacity The amount of capacity occupied by snapshots of volumes.

Spare Capacity Capacity that is reserved for system use, when recovery from failure is required. This capacity cannot be used for storage purposes.

Spare Percentage Policy This policy determines the amount of capacity that must always be reserved as free space.

Standby MDM An MDM node that is ready to use, with an ID, that has been locked to a specific ScaleIO system.

Storage Pool A sub-set of physical storage devices in a Protection Domain. Each storage device can only belong to one Storage Pool. User volumes will always use the storage of a single Storage Pool.

T

Thick Capacity Capacity allocated for thick volumes.

Thick Provisioned Volume In virtual storage, thick provisioning is a type of storage allocation in which the amount of storage capacity on a disk is pre-allocated on physical storage at the time the disk is created, meaning that the volume has all its capacity pre-allocated on creation.

Thin Capacity Capacity allocated for thin volumes.

Thin Provisioned Volume Thin provisioning is a method of optimizing the efficiency with which the available space is utilized in storage area networks (SAN). Thin provisioning operates by allocating disk storage space in a flexible manner among multiple users, based on the minimum space required by each user at any given time.

Throttling Throttling controls resource prioritization for rebuild and rebalance processes. Throttling can be controlled per Protection Domain or per Storage Pool (by configuring rebuild and rebalance policies).

Tie Breaker The Tie Breaker (TB) is an MDM that does not have a manager role, whose sole purpose is to help determine which MDM module is the manager that will become the master MDM and take control over the ScaleIO cluster.

The Tie Breaker ensures that there will always be one Master MDM achieving cluster quorum. In a 3-node cluster, there is one TB; in a 5-node cluster, there are two TBs.

U

Unavailable Capacity Capacity that is not being used, but is also unavailable (due to server outage).

Unused Capacity Capacity that is not currently being used for any purpose in the system.

V

Volume A general term referring to a storage device. In the ScaleIO system, a volume consists of multiple blocks spread evenly on Storage Pool devices.

W

Widget The full screen view can be minimized into a widget, which is a small window that floats on your screen, over other applications. Property sheets can also be minimized into widgets.

Write Misses Write requests that were not found in cache