

# Dell EMC ScaleIO

Version 2.x

## User Guide

P/N 302-003-286

REV 05

Copyright © 2016-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published February 2018

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

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

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

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

# 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:

<b>Bold</b>	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Used for full titles of publications referenced in text
<code>Monospace</code>	Used for: <ul style="list-style-type: none"><li>• System code</li><li>• System output, such as an error message or script</li><li>• Pathnames, filenames, prompts, and syntax</li><li>• Commands and options</li></ul>
<i>Monospace italic</i>	Used for variables
<code>Monospace bold</code>	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](mailto:techpubcomments@emc.com).

# CONTENTS

	<b>Preface</b>	<b>3</b>
	<b>Figures</b>	<b>13</b>
	<b>Tables</b>	<b>15</b>
<b>Part 1</b>	<b>Introduction</b>	<b>17</b>
<b>Chapter 1</b>	<b>Introduction to ScaleIO</b>	<b>19</b>
	What is ScaleIO?.....	20
	System requirements.....	20
	ScaleIO cluster components.....	21
	Physical server requirements.....	21
	Supported operating systems.....	22
	External SDC support.....	24
	GUI server requirements.....	24
	ScaleIO Gateway server requirements.....	25
	Other requirements.....	25
	New and changed features.....	26
	New features for ScaleIO.....	26
	Changed features.....	26
	Product limits.....	26
<b>Chapter 2</b>	<b>Architecture</b>	<b>29</b>
	ScaleIO Architecture Overview.....	30
	System.....	30
	Hardware.....	30
	Software.....	30
	The MDM cluster.....	31
	Storage definitions.....	33
	Protection Domains.....	33
	Storage Pools.....	33
	Fault Sets.....	35
	Naming.....	36
	Protection and load balancing.....	36
	Rebuild.....	36
	Rebuild throttling.....	37
	Rebalance.....	37
	Rebalance throttling.....	38
	Checksum protection.....	39
	ESX vStorage APIs for Array Integration (VAAI).....	39
	Caching.....	40
	Networking.....	43
	VMware limitation:.....	46
	Virtual IP Address.....	46
	Monitoring of SDC and SDS connections.....	47
	S.M.A.R.T. hardware monitoring.....	48

	List of approved RAID controllers.....	49
	Snapshots.....	50
	V-Trees.....	51
	Other functions.....	51
	Implementing ScaleIO.....	55
	Physical layer.....	55
	SAN virtualization layer.....	57
	Implementing ScaleIO over a virtual system.....	58
	Implementing ScaleIO in an ESXi-based system.....	58
	Xen implementation.....	61
	Maintenance.....	61
	Maintaining the physical layer.....	61
	Instant maintenance mode.....	62
	Maintaining the virtualization layer.....	62
	Management tools.....	63
	Configuring direct attached storage (DAS).....	63
<b>Part 2</b>	<b>Getting Started</b>	<b>65</b>
<b>Chapter 3</b>	<b>Licensing</b>	<b>67</b>
	Licensing overview.....	68
	Activating entitlements and installing a license file.....	69
	Activating an entitlement and downloading the license file.....	70
	Installing the license.....	73
	License file example.....	73
	Error messages.....	74
<b>Chapter 4</b>	<b>User Management</b>	<b>75</b>
	Overview.....	76
	User roles.....	76
	Logging in.....	77
	login.....	78
	logout.....	79
	Setting the User Authentication Method.....	79
	set_user_authentication_method.....	79
	Adding and modifying local users.....	80
	add_user.....	80
	delete_user.....	81
	modify_user.....	81
	query_users.....	82
	query_user.....	82
	reset_password.....	83
	set_password.....	83
	disable_admin.....	84
	Reset the admin user password.....	84
<b>Chapter 5</b>	<b>Creating and Mapping Volumes</b>	<b>87</b>
	Creating and mapping volumes overview.....	88
	Creating volumes.....	88
	Adding volumes.....	88
	Mapping a volume to an SDC.....	90
	Mapping volumes.....	90

<b>Part 3</b>	<b>Managing and Monitoring</b>	<b>93</b>
<b>Chapter 6</b>	<b>Managing System Objects</b>	<b>95</b>
	CLI basics.....	96
	Using SCLI in non-secure mode.....	97
	Syntax.....	97
	Getting help with the CLI.....	98
	Extend an existing ScaleIO system.....	99
	Managing the MDM cluster.....	99
	Replacing, or updating an IP address on a member of the MDM cluster.....	99
	Configure virtual IP addresses.....	106
	Managing SDC access to the MDM.....	106
	Add another IP address subnet to an MDM cluster.....	107
	Managing the SDSs and cache.....	109
	Modifying an SDS port during IO.....	109
	Managing Read Flash cache.....	110
	Managing read RAM cache.....	111
	SDC operations.....	111
	Updating the SDC driver with IP changes.....	112
	Detecting new volumes.....	112
	Query volumes using drv_cfg.....	113
	Query tgt objects using drv_cfg.....	113
	Query GUID using drv_cfg.....	114
	Query MDMs using drv_cfg.....	115
	Loading a configuration file using drv_cfg.....	115
	Adding an MDM using drv_cfg.....	116
	Modifying an MDM IP address using drv_cfg.....	117
	Permanent Device Loss state.....	119
	Managing ESX servers.....	120
	Modifying parameters on ESX servers.....	120
	Checking the SDC state on ESX servers.....	122
<b>Chapter 7</b>	<b>Security Management</b>	<b>123</b>
	Setting up SSH authentication on the ScaleIO Gateway.....	124
	Configuring SSL component authentication.....	124
	Internal component authentication.....	124
	External component authentication.....	124
	Workflow for self-signed security certificates.....	125
	Workflow for externally signed security certificates.....	125
	Using Keytool to add certificates to external components.....	126
	Managing SDC access to the MDM.....	128
	Approved encryption methods.....	129
	Login banner overview.....	129
	Setting up a login banner using the CLI.....	130
	Enabling or disabling preemptive acceptance of the login banner....	130
	Activating preemptive acceptance of the login banner.....	131
<b>Chapter 8</b>	<b>Opening the GUI and Logging In</b>	<b>133</b>
	Log in to the ScaleIO GUI.....	134
	Connection and disconnection information.....	134

<b>Chapter 9</b>	<b>GUI Features</b>	<b>135</b>
	GUI overview.....	136
	GUI conventions.....	136
	Alerts indicators.....	137
	Color codes.....	137
	Dashboard view.....	138
	Navigation tree.....	139
	System Settings and Logout menus.....	140
	Dimmer.....	140
	Widget.....	140
	Dashboard tiles.....	141
	Management (MDM) cluster status.....	144
	Frontend views.....	146
	Volumes.....	147
	SDCs.....	148
	Snapshots.....	148
	Adding Frontend objects to the filter.....	148
	Backend view.....	149
	Command menu (3).....	151
	Filter (1).....	151
	Table.....	151
	Alerts view.....	155
	Property Sheets.....	156
 <b>Chapter 10</b>	 <b>Monitoring the System using the GUI</b>	 <b>161</b>
	Viewing object properties in Backend and Frontend views.....	162
	Viewing licensing information.....	162
	Verifying your connection to the Management cluster.....	162
 <b>Chapter 11</b>	 <b>Configuring the System using the GUI</b>	 <b>163</b>
	Configuring capacity.....	164
	Adding SDSs and storage devices.....	164
	Removing SDSs and devices.....	167
	Adding, removing, and activating and inactivating capacity.....	168
	Activating devices.....	170
	Clearing device errors.....	170
	Setting device capacity limits.....	170
	Configuring cache.....	171
	Setting Read Flash Cache policy at Storage Pool level.....	171
	Setting Read Flash Cache policy at SDS level.....	172
	Adding Read Flash Cache devices.....	172
	Changing Read RAM Cache volume settings.....	173
	Removing Read Flash Cache devices.....	174
	Configuring Read RAM Cache (advanced, Backend).....	174
	Configuring volumes, volume trees, SDCs, and snapshots.....	176
	Adding volumes.....	176
	Restricted SDC mode.....	178
	Mapping and unmapping volumes.....	180
	Removing volumes.....	181
	Creating volume snapshots.....	182
	Removing snapshots.....	183
	Removing snapshots from a consistency group.....	185
	Increasing a volume's size.....	186
	Setting volume bandwidth and IOPS limits.....	187



	Entering and exiting SDS Instant Maintenance Mode.....	188
	Configuring Oscillating Failure counters.....	189
	Configuring Oscillating Failure counter parameters.....	189
	Resetting Oscillating Failure counters.....	192
	Viewing Oscillating Failure counters.....	193
	Applying Performance Profiles to system components.....	194
	Configuring I/O priorities and bandwidth use (advanced).....	195
	Application IOPS and bandwidth (advanced).....	195
	System IOPS and bandwidth (advanced).....	196
	Enabling and disabling Rebuild/Rebalance (advanced).....	197
	Using the background device scanner.....	197
	Enabling and disabling the background device scanner.....	198
	Resetting the background device scanner counters.....	199
	Modifying Checksum protection mode.....	199
	Renaming objects.....	200
	Approving pending security certificates.....	201
	Customizing system preferences.....	201
<b>Chapter 12</b>	<b>Using the VMware Plug-in</b>	<b>205</b>
	VMware Plug-in overview.....	206
	Configuring components.....	207
	Configuring components—basic.....	209
	Configuring components—advanced.....	210
	Viewing components.....	216
<b>Part 4</b>	<b>Reference</b>	<b>219</b>
<b>Chapter 13</b>	<b>Common Tasks</b>	<b>221</b>
	Install the ScaleIO GUI.....	222
	Log in to the ScaleIO GUI.....	222
	Connection and disconnection information.....	223
	Add LIA to a system to enable automated upgrade.....	223
	Associating ScaleIO volumes with physical disks.....	224
	Volume information - Linux.....	224
	Volume information - Windows.....	225
	Volume information - AIX.....	226
	Port usage and changing default ports.....	227
	Adding an external SDC to an existing ScaleIO system.....	228
	Installing SDC on an ESX server and connecting it to ScaleIO.....	228
	Installing SDC on a Linux server and connecting it to ScaleIO.....	229
	Install SDC on an AIX server and connect it to ScaleIO.....	230
	Installing SDC on a Windows server and connecting it to ScaleIO.....	230
	Changing the LIA configuration file.....	231
	Cleaning the ScaleIO VMware environment and performing a clean install.....	231
	Configuring ScaleIO devices in Linux LVM.....	233
	Configuring session timeout parameters.....	234
	Fixing keytool errors.....	234
	Error during rpm installation command.....	234
	Error during rpm upgrade command.....	234
	Installing Java on SUSE 12 servers.....	235
	SVM manual memory allocation.....	235
	Upgrading Java.....	237

	Mounting ScaleIO.....	237
	The ScaleIO Gateway web server isn't responding.....	238
	The ScaleIO Gateway (REST service, Installation Manager) may be disabled:.....	238
	The ScaleIO Gateway web server isn't responsive and the following error appears in the catalina log file:.....	239
	Upgrading the Gateway when a custom certificate is used.....	240
	Uploading a new OVA.....	240
	Using the same data network for different NICs.....	241
	What to do when the default self-signed certificate expires.....	241
	Add another IP address subnet to an MDM cluster.....	241
	Shutdown or restart a node gracefully.....	243
	Gracefully shut down or reboot a node.....	243
	Return the node to operation.....	245
	Deployment of ScaleIO using a non-root user.....	246
	Configure a non-root non-interactive sudo user.....	246
<b>Chapter 14</b>	<b>System events</b>	<b>249</b>
	System events overview.....	250
	Event format.....	250
	Viewing events.....	251
	Viewing events locally.....	251
	Event list.....	254
	Authentication.....	254
	CLI commands.....	255
	License and installation.....	256
	MDM.....	258
	SDC.....	260
	SDS.....	261
	Rebuild.....	264
<b>Chapter 15</b>	<b>ScaleIO on Xen</b>	<b>267</b>
	Overview of ScaleIO on Xen.....	268
	Adding a volume in XenServer environment.....	268
	Removing a ScaleIO volume from Xen.....	269
	Modifying the size of a ScaleIO volume.....	270
	Xen v6.5 High Availability.....	271
<b>Chapter 16</b>	<b>Configuring ScaleIO in OpenStack Environments</b>	<b>273</b>
	Overview.....	274
<b>Chapter 17</b>	<b>SNMP Trap Support</b>	<b>275</b>
	General.....	276
	Supported alerts and event numbering conventions.....	276
	CLASS/TYPE:.....	276
	ScaleIO Alerts in SNMP, GUI, REST, and ESRS.....	277
	Configure SNMP properties after deployment.....	303
	Create a Lockbox.....	303
	Configuring SNMP after deployment.....	304
	ScaleIO.mib file.....	305
<b>Chapter 18</b>	<b>ScaleIO SDC on AIX server</b>	<b>309</b>

SAN virtualization layer.....	310
SDC activities and monitoring.....	311
Enable error logging on AIX servers.....	313
Modify MDM IP address and GUID on AIX server.....	313

<b>Glossary</b>	<b>315</b>
-----------------	------------



# FIGURES

1	5-node MDM cluster.....	32
2	Protection Domains and Storage Pools.....	34
3	Protection Domains, Storage Pools, and Fault Sets.....	35
4	ScaleIO system deployed on a single network.....	45
5	ScaleIO system deployed on separate networks.....	46
6	Snapshot operations.....	50
7	V-Tree diagram.....	51
8	Physical layout example—3-node cluster.....	55
9	Physical layout example—3-node cluster.....	56
10	ScaleIO implementation on ESX.....	59
11	ScaleIO Xen virtual machine architecture.....	61
12	Licensing LAC email.....	69
13	License file example.....	73
14	Add Volume window.....	90
15	Map Volumes window after mapping is complete.....	91
16	Alerts indicators.....	137
17	Dashboard controls.....	139
18	Widget, showing Return to Full-Screen button.....	141
19	Dashboard tiles.....	142
20	Frontend > Volumes view.....	147
21	Frontend > SDCs view.....	148
22	Frontend > filtered Snapshots view .....	148
23	Backend view.....	150
24	Backend filter.....	151
25	Alerts view.....	155
26	Multiple floating Property Sheets.....	158
27	Example of a Property Sheet for an SDS.....	159
28	Add SDS window.....	166
29	Add Device window showing command validation.....	167
30	Set Volume Use Read RAM Cache window.....	174
31	RAM Read Cache configuration at Storage Pool level.....	175
32	RAM Read Cache configuration at SDS level.....	176
33	Add Volume window.....	178
34	Map Volumes window after mapping is complete.....	181
35	Unmap Volumes window.....	181
36	Remove Volumes window.....	182
37	Snapshot Volume window.....	183
38	Removal of a volume and associated snapshots.....	185
39	Removal of snapshots of a specific volume.....	185
40	Remove Consistency Groups window.....	186
41	Increase Volumes' Size window.....	187
42	Set Volume Limits window.....	188
43	Configure Oscillating Failure counters—System.....	191
44	Configure Oscillating Failure counters—Protection Domain or Storage Pool.....	191
45	Reset Oscillating Failure counters—System.....	192
46	Configure Oscillating Failure counters—Protection Domain or Storage Pool.....	193
47	Set Performance Profile window.....	195
48	Set Network Throttling window.....	197
49	Background device scanner configuration.....	199
50	Reset Background Device Scanner Counters window.....	199
51	Configure Use Checksum window.....	200
52	User Preferences window.....	202
53	Configure virtual IPs dialog box.....	216

## FIGURES

# TABLES

1	Server physical requirements.....	21
2	Supported operating systems - ScaleIO components.....	23
3	Product limits.....	26
4	MDM cluster modes.....	32
5	Caching modes.....	40
6	Caching support matrix.....	42
7	IP address configurations in ScaleIO (based on CSV file).....	44
8	eLicensing terminology(continued).....	69
9	Licensing error messages.....	74
10	Local and LDAP user roles and permissions.....	76
11	GUI task overview.....	136
12	GUI color codes.....	137
13	Alert symbols and color codes.....	138
14	Dashboard tiles.....	142
15	Management node icons with normal operational status (green).....	145
16	Management node status indications and color codes.....	146
17	Table view options.....	152
18	User Preferences.....	202
19	Plug-in activity matrix.....	207
20	Default ports.....	227
21	ScaleIO Alerts in SNMP, GUI, REST, and ESRS.....	277
22	SDC activities and monitoring.....	311





# PART 1

## Introduction

This section describes an overview of the benefits and system requirements of ScaleIO.

Chapters include:

[Chapter 1, "Introduction to ScaleIO"](#)

[Chapter 2, "Architecture"](#)



# CHAPTER 1

## Introduction to ScaleIO

This section introduces ScaleIO.

- [What is ScaleIO?](#) ..... 20
- [System requirements](#)..... 20
- [New and changed features](#).....26
- [Product limits](#).....26

## 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 ESXi and RHEL operating systems are currently supported. For more information, see the operating system support tables.

The 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.

---

### Note

In the documentation set, there are references that are specific to either ESXi or RHEL operating systems, but not to both (for example, vCenter). These differences are not marked in most places.

---

## 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 RCache 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: <ul style="list-style-type: none"> <li>• Intel or AMD x86 64-bit (recommended)</li> <li>• Intel or AMD x86 32-bit (for Xen only)</li> </ul>
Physical memory	ScaleIO component requirements: <ul style="list-style-type: none"> <li>• 500 MB RAM for the Meta Data Manager (MDM)</li> <li>• 500 MB RAM for each ScaleIO Data Server (SDS)</li> <li>• 50 MB RAM for each ScaleIO Data Client (SDC)</li> </ul> DAS Cache memory requirements (ScaleIO Ready Node, non-XenServers only). Add to every SVM/node that will be using DAS Cache: <ul style="list-style-type: none"> <li>• 1U1N servers—500 MB RAM</li> </ul>

**Table 1** Server physical requirements (continued)

Component	Requirement
	<ul style="list-style-type: none"> <li>2U1N servers—1 GB RAM</li> </ul> <p>To calculate SVM memory allocation, use the formulas provided in the <i>ScaleIO Deployment Guide</i>.</p>
Disk space	<ul style="list-style-type: none"> <li>1 GB for each physical node or Xen hypervisor</li> <li>10 GB for VMware topologies</li> </ul>
Connectivity	<p>One of the following:</p> <ul style="list-style-type: none"> <li>1 GbE or 10 GbE (recommended) network</li> <li>IP-over-InfiniBand network</li> </ul> <p>Dual-port network interface cards (recommended)</p> <p>Ensure the following:</p> <ul style="list-style-type: none"> <li>There is network connectivity between all components.</li> <li>Network bandwidth and latency between all nodes is acceptable, according to application demands.</li> <li>Ethernet switch supports the bandwidth between network nodes.</li> <li>MTU settings are consistent across all servers and switches.</li> <li>The following TCP ports are not used by any other application, and are open in the local firewall of the server: <ul style="list-style-type: none"> <li>MDM: 6611 and 9011</li> <li>Tie Breaker: 9011</li> <li>SDS: 7072. Multiple SDS (not supported on Windows): 7073-7076</li> <li>Light Installation Agent (LIA): 9099</li> <li>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).</li> </ul> </li> <li>The following UDP port is open in the local firewall of the server: <ul style="list-style-type: none"> <li>SNMP traps: 162</li> </ul> </li> </ul> <hr/> <p><b>Note</b></p> <p>You can change the default ports. For more information, see “Changing default ports” in the user documentation.</p> <hr/>

## Supported operating systems

The following is a 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/elhome>.

**Table 2** Supported operating systems - ScaleIO components

Operating system	Requirement
Linux	<p>Supported versions:</p> <ul style="list-style-type: none"> <li>CentOS 6.x-7.x, Oracle Linux 6.5/7.x</li> <li>Red Hat 6.x-7.x</li> <li>SUSE 11.3, 11.4, 12, 12.1, 12.2</li> <li>Ubuntu 14.04, Ubuntu 16.04</li> </ul> <hr/> <p><b>Note</b></p> <p>Before deploying SDC or RFCache on Ubuntu servers, you must prepare the environment, as described in the <i>EMC ScaleIO Deployment Guide</i>.</p> <hr/> <p>Packages required for all components, all Linux flavors:</p> <ul style="list-style-type: none"> <li>numactl</li> <li>libaio</li> </ul> <p>Additional packages required for MDM components:</p> <ul style="list-style-type: none"> <li>bash-completion (for SCLI completion)</li> <li>Latest version of Python 2.X</li> </ul> <p>When installing the MDM component on Linux CentOS 6 or RHEL 6 hosts (for software-only systems), set the shared memory parameter in the <code>/etc/sysctl.conf</code> file to at least the following value: <code>kernel.shmmax=209715200</code>. To use this value, type the <code>sysctl -p</code> command.</p> <p>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.</p> <p>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:</p> <ul style="list-style-type: none"> <li><code>libopenssl1_0_0-1.0.1g-0.40.1.x86_64.rpm</code></li> <li><code>openssl1-1.0.1g-0.40.1.x86_64.rpm</code></li> </ul> <p>To use LDAP, ensure that OpenLDAP 2.4 is installed on all servers.</p>
Windows	<p>Supported versions:</p> <ul style="list-style-type: none"> <li>2008 R2, 2012, 2012 R2, or 2016 (in v2.0.1.1 and later). Server Core editions are not supported. (For ScaleIO Ready Node, 2008 R2 and 2012 are not supported.)</li> <li>For VxRack Node 100 Series, only 2012 R2 is supported.</li> <li>On all MDM servers, install the EMC-provided <code>PythonModulesInstall.exe</code> 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 <a href="https://support.emc.com">https://support.emc.com</a>.</li> <li>To install SDC or RFCache on 2008 R2, ensure that Microsoft Security Update KB3033929 is installed.</li> </ul> <p>To use the secure authentication mode, ensure that these are installed on all servers in the system:</p> <ul style="list-style-type: none"> <li>OpenSSL 64-bit v1.0.1 or later (v1.1, however, is not supported)</li> </ul>

**Table 2** Supported operating systems - ScaleIO components (continued)

Operating system	Requirement
	<ul style="list-style-type: none"> <li>Visual C++ redistributable 2010 package, 64-bit</li> </ul> <p>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.</p>
Hypervisors	<ul style="list-style-type: none"> <li>VMware ESXi OS: 5.5 U3, 6.0 U3, or 6.5, managed by vCenter 5.5, 6.0, or 6.5</li> <li>Hyper-V</li> <li>XenServer 6.5 or 7.0</li> </ul> <hr/> <p><b>Note</b></p> <p>OpenSSL 64-bit v1.0.1 is supported on XenServer 6.5 SP1 (or later)</p> <hr/> <ul style="list-style-type: none"> <li>Red Hat KVM</li> </ul>

## External SDC support

In addition to being supported on all ScaleIO operating systems, SDC can be deployed on external servers.

Component	Requirement
Supported external servers	<ul style="list-style-type: none"> <li>UNIX: AIX 7.2</li> <li>hLinux: 4.x/5.x</li> </ul>

## GUI server requirements

Component	Requirement
Supported operating systems	<ul style="list-style-type: none"> <li>Windows: <ul style="list-style-type: none"> <li>7, 2008 R2, 10, 2012 or 2012 R2, 2016.</li> </ul> Server Core editions are not supported. </li> <li>Linux: <ul style="list-style-type: none"> <li>CentOS 6.x-7.x, Oracle Linux 6.5/7.x</li> <li>Red Hat 6.x-7.x</li> <li>SUSE 11.3, 11.4, 12, 12.1, 12.2</li> <li>Ubuntu 14.04, Ubuntu 16.04</li> </ul> </li> </ul>
Other	<ul style="list-style-type: none"> <li>v1.8 (64-bit), build 149 or earlier.</li> </ul> <p>You can download previous versions from this link: <a href="http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html">http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html</a></p> <p>For SUSE 12, see “Installing Java on SUSE 12 servers.” in the <i>ScaleIO Deployment Guide</i> or the <i>ScaleIO User Guide</i>.)</p>



Component	Requirement
	<ul style="list-style-type: none"> <li>Screen resolution: 1366 x 768 minimum</li> </ul>

## ScaleIO Gateway server requirements

Component	Requirement
Supported operating systems	<ul style="list-style-type: none"> <li>Windows 2008 R2, 2012 R2, or 2016, including the Visual C++ redistributable 2010 package, 64-bit. Server Core editions are not supported.</li> <li>Linux: <ul style="list-style-type: none"> <li>CentOS 6.x-7.x</li> <li>Oracle Linux 6.5/7.x</li> <li>Red Hat 6.x-7.x</li> <li>SUSE 11.3, 12, 12.1, and 12.2</li> <li>Ubuntu 14.04, Ubuntu 16.04</li> </ul> </li> </ul> <p>Every server requires 2 cores and a minimum of 3 GB available RAM.</p>
Connectivity	<p>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).</p> <p>You can change the default ports. For more information, see “Changing default ports” in the user documentation.</p>
Supported web browsers	<ul style="list-style-type: none"> <li>Internet Explorer 10, or later</li> <li>Firefox, version 42, or later</li> <li>Chrome, version 45, or later</li> </ul>
Java	<ul style="list-style-type: none"> <li>v1.8 (64-bit), build 149 or earlier.</li> </ul> <p>You can download previous versions from this link: <a href="http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html">http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html</a></p>
Other	<ul style="list-style-type: none"> <li>For a Windows Gateway, the Windows Management Instrumentation service must be enabled on the IM server and on all Windows ScaleIO nodes.</li> <li>Do not install the Gateway on a server on which RFCache will be enabled or on which SDC will be installed.</li> <li>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.</li> </ul>

## 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*.

---

## New and changed features

### New features for ScaleIO

Learn about the new features introduced in the ScaleIO 2.5-<build> software.

#### Deployment with non-root user

In Linux environments, you can now deploy a ScaleIO cluster using a non-root sudo user in non-interactive mode.

#### SDC IP address association

To enable you to have more control over your ScaleIO system, you can set your system to run in restricted mode. This mode requires you to map volumes only to SDCs which have been previously approved by the user, by configuring them either by IP address or GUID.

#### Multiple LDAP server support

ScaleIO now supports the use of multiple LDAP servers.

#### SDC disconnection alerts

The system now generates alerts when disconnections occur between SDCs and SDSs.

#### Operating system support

ScaleIO now also supports the following operating systems:

- AIX 7.2 (for SDC core component only)
- ESXi 6.5 U1
- XEN 7.1
- XEN 7.2

### Changed features

Learn about the changed features for ScaleIO 2.5-<build>.

## 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

**Table 3** Product limits (continued)

Item	Limit
Minimum Storage Pool capacity	300 GB
Volume size	8 GB—1 PB
Maximum number of volumes/snapshots in system	32,768 <sup>a</sup>
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 <sup>a</sup>
Maximum devices (disks) per SDS server	64 <sup>b</sup>
Maximum devices (disks) per Storage Pool	300 <sup>a</sup>
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. <sup>c</sup>
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

The following topics describe the ScaleIO Ready Node architecture.

• <a href="#">ScaleIO Architecture Overview</a> .....	30
• <a href="#">System</a> .....	30
• <a href="#">The MDM cluster</a> .....	31
• <a href="#">Storage definitions</a> .....	33
• <a href="#">Protection and load balancing</a> .....	36
• <a href="#">ESX vStorage APIs for Array Integration (VAAI)</a> .....	39
• <a href="#">Caching</a> .....	40
• <a href="#">Networking</a> .....	43
• <a href="#">Virtual IP Address</a> .....	46
• <a href="#">Monitoring of SDC and SDS connections</a> .....	47
• <a href="#">S.M.A.R.T. hardware monitoring</a> .....	48
• <a href="#">Snapshots</a> .....	50
• <a href="#">Other functions</a> .....	51
• <a href="#">Implementing ScaleIO</a> .....	55
• <a href="#">Implementing ScaleIO over a virtual system</a> .....	58
• <a href="#">Maintenance</a> .....	61
• <a href="#">Management tools</a> .....	63
• <a href="#">Configuring direct attached storage (DAS)</a> .....	63

# 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 that reside 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)	<ul style="list-style-type: none"> <li>• Master MDM</li> <li>• Slave MDM</li> <li>• Tie Breaker</li> </ul>	3-node cluster has two copies of the repository, thus can withstand one MDM cluster failure.
5-node	<ul style="list-style-type: none"> <li>• Master MDM</li> <li>• Two Slave MDM</li> <li>• Two Tie Breaker</li> </ul>	5-node cluster has three copies of the repository, thus can withstand two MDM cluster failure.
Single-node	<ul style="list-style-type: none"> <li>• Master MDM</li> </ul>	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 MDM cluster IP address limit is 16 IP addresses, which includes all cluster members (Master, Slave, Standby Master, and Standby Slaves).

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 available 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 35 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
- vSphere plug-in

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.

---

### Note

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

---

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

## 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.

---

### 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.

---

- 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.

## 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 vaa1 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 and NVMe 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 only

---

#### Note

DAS Cache is not supported on RHEL 7.4 operating systems nor on PowerEdge 14G 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.	<p>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.</p> <hr/> <p><b>Note</b></p> <p>The amount that may be allocated is limited, and can never be the maximum available RAM.</p> <hr/>	Disabled, except when storage-only nodes are deployed.



**Table 5** Caching modes (continued)

Mode	Description	Considerations	Default Setting
Read Flash Cache (RFcache)	Read-only caching performed by one or more dedicated SSD or NVMe SSD devices in the server.	<p>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.</p> <p>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.</p> <p>The RFcache driver must be installed during deployment. Caching devices can be defined either during the installation process or after deployment.</p> <p>Limitations:</p> <p>RFcache does not support partitions on devices installed on Windows nodes.</p> <p>Support matrix:</p> <ul style="list-style-type: none"> <li>• An RFcache device (flash device) can be partitioned only on Linux.</li> <li>• An SDS storage/source device cannot be partitioned if it needs to be accelerated by RFcache.</li> <li>• An SDS storage/source device as a file (over file system), cannot be accelerated by RFcache.</li> </ul>	
CacheCade	Read and write-back caching performed by one or more dedicated SSD devices in the server.	<p>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.</p> <p>Two SSD devices can be allocated to a shared cache pool, up to a maximum size of 512 GB in total.</p>	Disabled

**Table 5** Caching modes (continued)

Mode	Description	Considerations	Default Setting
		<p><b>Note</b></p> <p>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.</p>	
DAS Cache	Read and write-back caching performed by one or more dedicated SSD devices in the server	<p>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.</p> <p><b>Note</b></p> <p>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.</p>	Disabled

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 PowerEdge 13G servers	Yes	Yes	Yes	
ScaleIO Ready Node PowerEdge 14G servers	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*.

---

## Note

The MDM cluster IP address limit is 16 IP addresses, which includes all cluster members (Master, Slave, Standby Master, and Standby Slaves).

---

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.

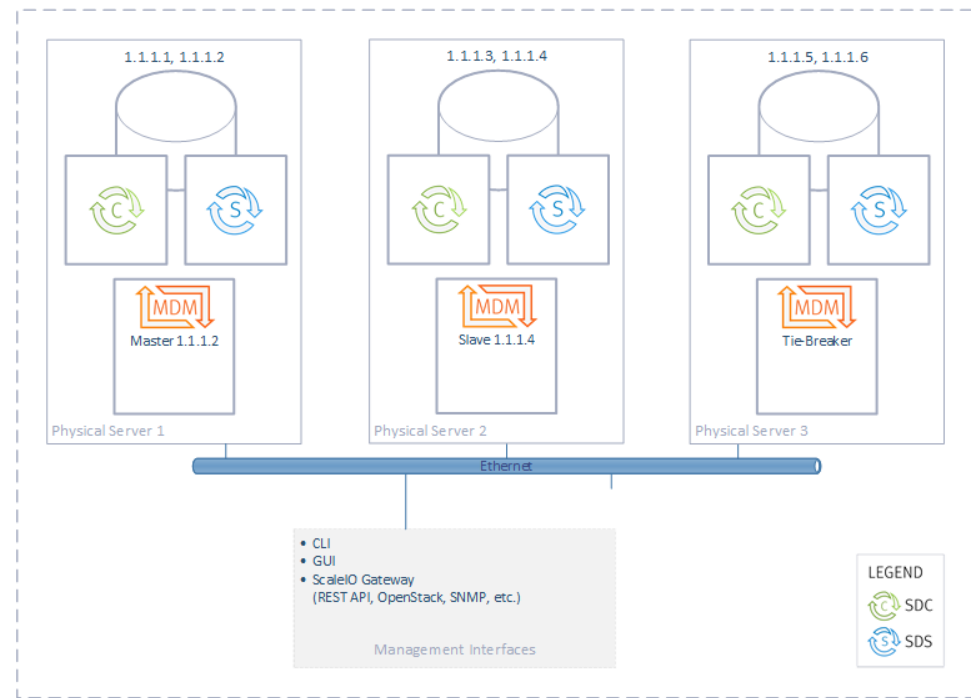
---

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:

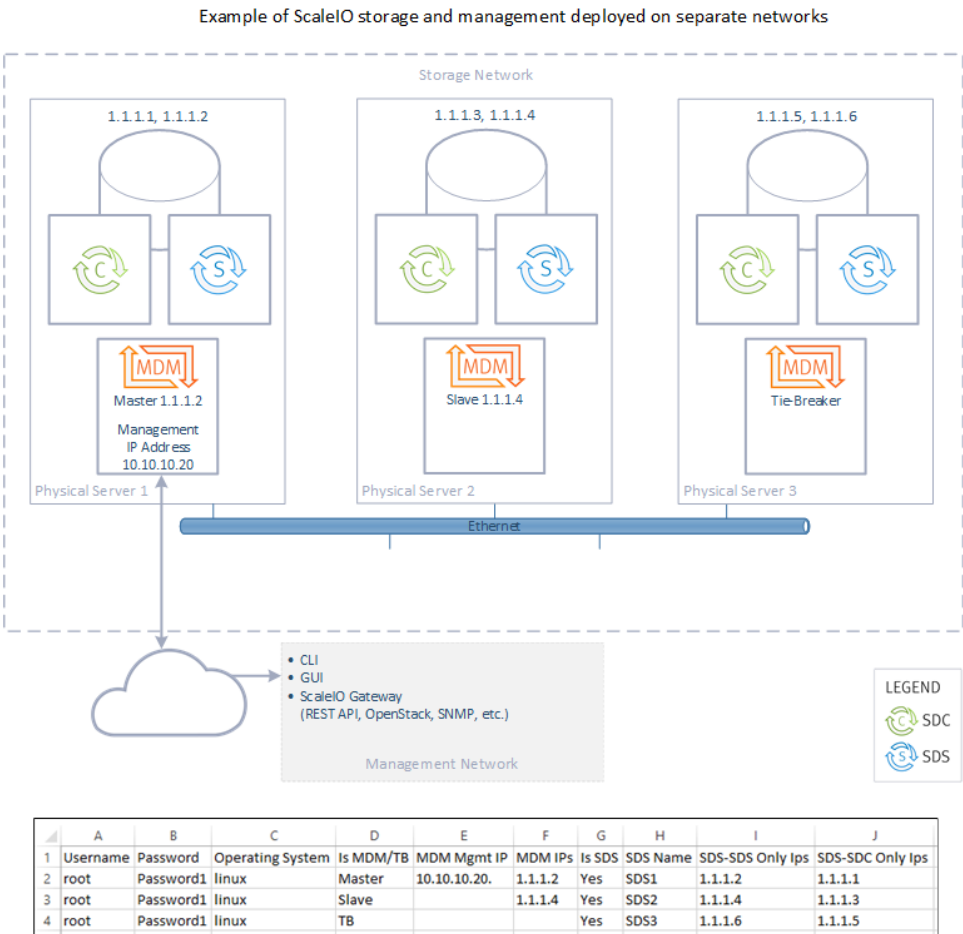
**Figure 4** ScaleIO system deployed on a single network

Example of a ScaleIO system deployed on a single network



	A	B	C	D	E	F	G	H	I	J
1	Username	Password	Operating System	Is MDM/TB	MDM Mgmt IP	MDM IPs	Is SDS	SDS Name	SDS-SDS Only Ips	SDS-SDC Only Ips
2	root	Password1	linux	Master		1.1.1.2	Yes	SDS1	1.1.1.2	1.1.1.1
3	root	Password1	linux	Slave		1.1.1.4	Yes	SDS2	1.1.1.4	1.1.1.3
4	root	Password1	linux	TB			Yes	SDS3	1.1.1.6	1.1.1.5

**Figure 5** ScaleIO system deployed on separate networks



## 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](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

Virtual IP addresses can be defined for the MDM cluster.

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.

## Monitoring of SDC and SDS connections

The system monitors all connections between SDCs and SDSs and sends out an alert when an active connection between an SDC and an SDS goes down.

To effectively monitor SDC and SDS connections, the MDM collects connectivity updates from all of the SDCs. The MDM posts events whenever an SDC connects to or disconnects from a specific SDS IP address. The MDM frequently analyzes the connectivity status to determine the current system state. The system does not send out alerts for temporary connectivity issues that are resolved in less than 10 seconds.

The following are the possible connectivity states between SDCs and SDSs in the system:

- All connected
- One SDC is disconnected from one SDS
- One SDC is disconnected from one SDS IP address
- One SDC is disconnected from all SDSs
- All SDCs are disconnected from one SDS
- All SDCs are disconnected from one SDS IP address
- All SDCs are disconnected from all SDSs
- Multiple disconnections

When the system's connectivity state changes to any state other than `All Connected`, an alert is displayed in the GUI and is written to the MDM event log. Once an alert is generated, you can use the SCLI to query details on the disconnection using the command `scli --query_sdc_to_sds_disconnections`. For more information about running SCLI commands, see the *CLI Reference Guide*.

The MDM does not monitor the connectivity state of SDCs or SDSs in the following scenarios:

- SDS is in maintenance mode
- SDS is disconnected from the MDM
- SDS is in the process of being removed
- SDC is disconnected from the MDM for more than two minutes
- SDC is not approved

The alerts are detailed in "Alerts in SNMP, GUI, REST, and ESRS" in the *ScaleIO User Guide*.

## 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 percccli 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 REST API Reference Guide*.

You can use the GUI to monitor S.M.A.R.T.-related alerts in the **Alerts** view. Configuration and status information are shown in the **Backend > Physical > S.M.A.R.T.** 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	<ul style="list-style-type: none"> <li>Model Name: Smart Array P440ar</li> <li>Vendor Name: HP</li> <li>Firmware Version: 3.56</li> <li>Driver Version: 3.4.10</li> <li>Driver Name: hpsa</li> <li>PCI Address: 0000:03:00.0</li> </ul>
DELL	<ul style="list-style-type: none"> <li>Model Name: PERC H730 Mini</li> <li>Vendor Name: Dell</li> <li>Firmware Version: 25.3.0.0016</li> <li>Driver Version: 06.807.10.00-rh</li> <li>Driver Name: megaraid_sas</li> <li>PCI Address: 00:02:00:00</li> </ul>
LSI	<ul style="list-style-type: none"> <li>Model Name: LSI MegaRAID SAS 9271-8i</li> <li>Vendor Name: LSI</li> <li>Firmware Version: 23.12.0-0021</li> <li>Driver Version: 06.810.09.00-rh</li> <li>Driver Name: megaraid_sas</li> <li>PCI Address: 00:82:00:00</li> </ul>
	<ul style="list-style-type: none"> <li>Model Name: LSI MegaRAID SAS 9271-8i</li> <li>Vendor Name: LSI</li> <li>Firmware Version: 23.12.0-0018</li> <li>Driver Version: 06.805.06.01-rc</li> <li>Driver Name: megaraid_sas</li> <li>PCI Address: 00:82:00:00</li> </ul>
	<ul style="list-style-type: none"> <li>Model Name: LSI MegaRAID SAS 9271-8i</li> <li>Vendor Name: LSI</li> <li>Firmware Version: 23.12.0-0021</li> <li>Driver Version: 06.805.06.01-rc</li> <li>Driver Name: megaraid_sas</li> <li>PCI Address: 00:82:00:00</li> </ul>

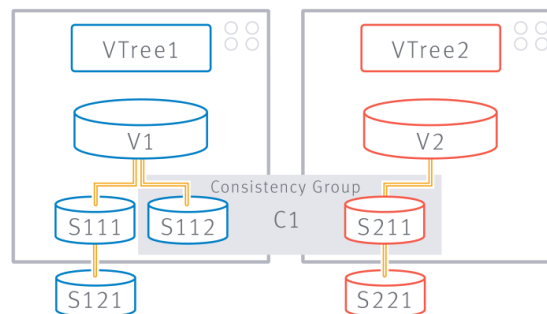
Manufacturer	Specifications
	<ul style="list-style-type: none"> <li>Model Name: LSI MegaRAID SAS 9271-8i</li> <li>Vendor Name: LSI</li> <li>Firmware Version: 23.34.0-0005</li> <li>Driver Version: 06.810.08.00</li> <li>Driver Name: megaraid_sas</li> <li>PCI Address: 00:82:00:00</li> </ul>

## 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 V-Tree (short for volume tree). 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. In [Figure 6](#) on page 50, in RED, S211 is a snapshot of V2. Since S112 and S211 were taken together, they compose a consistency group 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

Oscillating failure counters	Description
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)
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.

**Figure 8** Physical layout example—3-node cluster



## 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 9** 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 =  $5 * (\text{gross capacity} - \text{used capacity})$   
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

- When a volume is defined on an AIX SDC, one device is created with the following pathnames:
  - A block device, named /dev/scinidX...n, where *X* is a number, starting from “0.”
  - A raw device, named /dev/rscinidX...n, where *X* is a number, starting from “0.”

In general, mapping SDCs to AIX raw devices will yield best performance. If you are using the device to create a filesystem, use the block device.

- 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.

### Implementing ScaleIO in an ESXi-based system

#### 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:



- 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 35.

### 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 11** 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 I/Os; 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.

Device data is erased when devices are added to SDS. When adding a device to an SDS, ScaleIO will check that the device is clear before adding it. An error will be returned, per device, if it is found not to be clear. You can override this check by using the force device takeover option.

The following devices are considered to be not "clear," and thus cannot be added to SDS:

- **Linux** - A complete device with either a filesystem or partition, or a partitioned device with a filesystem.
- **Windows** - A complete device with a partition, or a partitioned device with a filesystem.
- **ESX** - Same as above, depending on the OS of the SVM where the SDS is installed.

**Limitations:**

- SAN devices will not be prevented from being added.
- Devices in an LVM group cannot be added to an SDS.
- Within the database devices, only Oracle ASM devices can be detected and blocked.

**NOTICE**

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 "Preparing devices on Windows servers" section in the *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 following chapters 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

The following topics describe how to obtain and activate the electronic license for your ScaleIO software.

- [Licensing overview](#) ..... 68
- [Activating entitlements and installing a license file](#) ..... 69
- [License file example](#) ..... 73
- [Error messages](#) ..... 74

## Licensing 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, <a href="http://support.emc.com">http://support.emc.com</a> , 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 12** 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://\[REDACTED\]](https://[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:

SOFTWARE LICENSING CENTRAL ACTIVATE ENTITLEMENTS LICENSES USAGE REPORTS HELP

Home » Activate »

## ACTIVATE

[Start Over](#)

- ✓ SELECT PRODUCTS
- ✓ COMPANY DETAILS
- 3 SELECT A MACHINE**
- 4 ENTER DETAILS
- 5 REVIEW
- 6 COMPLETE

### STEP 3: SELECT A MACHINE

#### SEARCH MACHINES [Search Tips](#)

% = supports partial search criteria

% Machine Name

% Locking ID

[Advanced Search](#)

**SEARCH**

#### ADD A NEW MACHINE

Don't see the machine you need?  
Add a new machine here.

Machine Name

**SAVE MACHINE & CONTINUE TO NEXT STEP**

#### MACHINE RESULTS

1-30 of 40 machines

Machine Name	Locking IDs	Product Lines

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:

Home » Activate »

## ACTIVATE

[Start Over](#)

- ✓ SELECT PRODUCTS
- ✓ COMPANY DETAILS
- ✓ SELECT A MACHINE
- 4 ENTER DETAILS**
- 5 REVIEW
- 6 COMPLETE

### STEP 4: ENTER PRODUCT QUANTITIES & MACHINE DETAILS

Products	Installed	Available	Quantity to Activate <a href="#">?</a>
<b>PRODUCT LINE: ScaleIO</b>			
EMC SIO SW lic key delivery=CB Product # 436-110-229	0	15	<b>5</b>

#### ENTER MACHINE DETAILS [Machine Details FAQ](#)

\* = field is required

Machine Name: SIO\_TEST2 | [Change Machine](#) [?](#)

\*ScaleIO Installation ID: 123456789abcdef | [X](#)

[BACK](#)
[X CANCEL](#)

[NEXT: REVIEW >](#)

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 13** License file example

<div style="border-left: 1px solid black; padding-left: 5px; margin-bottom: 10px;">Header</div> <div style="border-left: 1px solid black; padding-left: 5px; margin-bottom: 10px;">Base capacity</div> <div style="border-left: 1px solid black; padding-left: 5px;">Enterprise features</div>	<pre>##### # EMC License File # Activation Date: Apr 10, 2014 08:49:41 AM # Activated By: robert grosso # Type:UNSERVED ##### INCREMENT SIO_BASE EMCLM 1.0 permanent uncunted \   VENDOR_STRING=CAPACITY=200;CAPACITY_UNIT=TB;installation_id=;SWID=SL6LC2B6HB4YSK;PLC=SIO; \   HOSTID=ANY dist_info="ACTIVATED TO SIOTEST CO." ISSUER=EMC \   ISSUED=10-Apr-2014 NOTICE="ACTIVATED TO License Site Number: \   PTA10APR20141109559" SN=2169155 SIGN="004A FD6C 87EC 2F63 248F \   FE9B A852 C700 8608 8332 21F2 9C72 5744 759C D6FE" [INCREMENT SIO_SNAPSHOTS EMCLM 1.0 permanent uncunted \   VENDOR_STRING=CAPACITY=200;CAPACITY_UNIT=TB;installation_id=;SWID=SL6LC2B6HB4YSK;PLC=SIO; \   HOSTID=ANY dist_info="ACTIVATED TO SIOTEST CO." ISSUER=EMC \   ISSUED=10-Apr-2014 NOTICE="ACTIVATED TO License Site Number: \   PTA10APR20141109559" SN=2169155 SIGN="004C 6427 4589 3BC4 2656 \   EC20 3A51 6300 6466 F7A7 566A 59AD 088D 2BEA DB10" INCREMENT SIO_QoS EMCLM 1.0 permanent uncunted \   VENDOR_STRING=CAPACITY=200;CAPACITY_UNIT=TB;installation_id=;SWID=SL6LC2B6HB4YSK;PLC=SIO; \   HOSTID=ANY dist_info="ACTIVATED TO SIOTEST CO." ISSUER=EMC \   ISSUED=10-Apr-2014 NOTICE="ACTIVATED TO License Site Number: \   PTA10APR20141109559" SN=2169155 SIGN="0069 FC1E 0AF5 FEFF F20C \   EBB2 146A 7A00 3C7E 42B0 6D0F 7D78 4560 983C E534" INCREMENT SIO_OBFUSCATION EMCLM 1.0 permanent uncunted \   VENDOR_STRING=CAPACITY=200;CAPACITY_UNIT=TB;installation_id=;SWID=SL6LC2B6HB4YSK;PLC=SIO; \   HOSTID=ANY dist_info="ACTIVATED TO SIOTEST CO." ISSUER=EMC \   ISSUED=10-Apr-2014 NOTICE="ACTIVATED TO License Site Number: \   PTA10APR20141109559" SN=2169155 SIGN="00D7 202C 9FCA E6F0 F08D \   26A0 BB01 9700 E681 2976 4892 BFC7 B96A 229B 73C6"</pre>
--	---

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

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 the Installation ID was entered 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

The following topics describe how to create and manage users.

• <a href="#">Overview</a> .....	76
• <a href="#">User roles</a> .....	76
• <a href="#">Logging in</a> .....	77
• <a href="#">Setting the User Authentication Method</a> .....	79
• <a href="#">Adding and modifying local users</a> .....	80

## 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 the *ScaleIO REST API Reference Guide*.

- 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—with the exception of 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 the 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 Backend operations only (Protection Domains,		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
			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
--password password1
```

You will be prompted to enter the password.

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

```
scli --mdm_ip 10.10.10.3,10.10.10.4 --login --username
JohnDoe@ldap.acme.com --password password1 --ldap_authentication
```

The default user created during setup is the SuperUser, with the *admin* username.

## 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

#### **--accept\_banner\_by\_scripts\_only**

Preemptive approval of login banner

### 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, enter "q" to quit the login banner, and then enter "y" to 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 and LDAP user 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 the administrator role can manage system users, including adding new users and deleting existing users, modifying user credentials, and resetting user passwords.

The following CLI commands allow you to manage 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.

Each user name should conform to the following rules:

1. Contains fewer than 32 characters
2. Contains only alphanumeric and punctuation characters (when punctuation characters are being used, you may need to use the " or ' characters in order to allow it).
3. Is unique within the object type

**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

**--new\_password <NEW\_PASSWORD>**

User's new password

---

#### Note

In Linux, to prevent the password from being recorded in the history log, omit the `old_password` or `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 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 Superuser.

The Superuser 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 Superuser in order to ensure that all users are associated with specific user roles.

**Note**

To re-enable the Superuser, 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
```

**Reset the admin user password**

You can reset the password of the default admin user (Supeuser) using the combination of a file written to the MDM and the `reset_admin` CLI command.

**Before you begin**

Ensure that you are using the admin user with Superuser permissions.

**Note**

The procedure refers only to the default admin user with Superuser permissions, which was created during the system setup.

**Procedure**

1. Create a text file named `MDM_SERVICE_MODE` on the MDM in the location corresponding to your operating system:

- **Windows:** C:\Program Files\emc\scaleio\MDM\logs\MDM\_SERVICE\_MODE.txt
- **Linux:** /opt/emc/scaleio/mdm/logs/MDM\_SERVICE\_MODE.txt

2. In the body of the file, type the text **Reset Admin**, and save the file.
3. From the CLI, run the **reset\_admin** command:

```
scli --reset_admin
```

## Results

The admin user password is reset to **admin**.

## reset\_admin

Reset the default Superuser.

Reset the password of the default admin user with Superuser permissions.

**reset\_admin**

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

## Syntax

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

## Parameters

**--i\_am\_sure**

Skip the safety questions for command execution.

## Example

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



# CHAPTER 5

## Creating and Mapping Volumes

The following topics describe 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.

- [Creating and mapping volumes overview](#)..... 88
- [Creating volumes](#)..... 88
- [Mapping a volume to an SDC](#)..... 90

## 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.

#### Before you begin

There must be at least three SDS nodes in the system and there must be sufficient capacity available.

---

#### Note

For the minimum size of an SDS, see [System requirements](#) on page 20.

---

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 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 Read RAM Cache volume settings](#) on page 173.

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

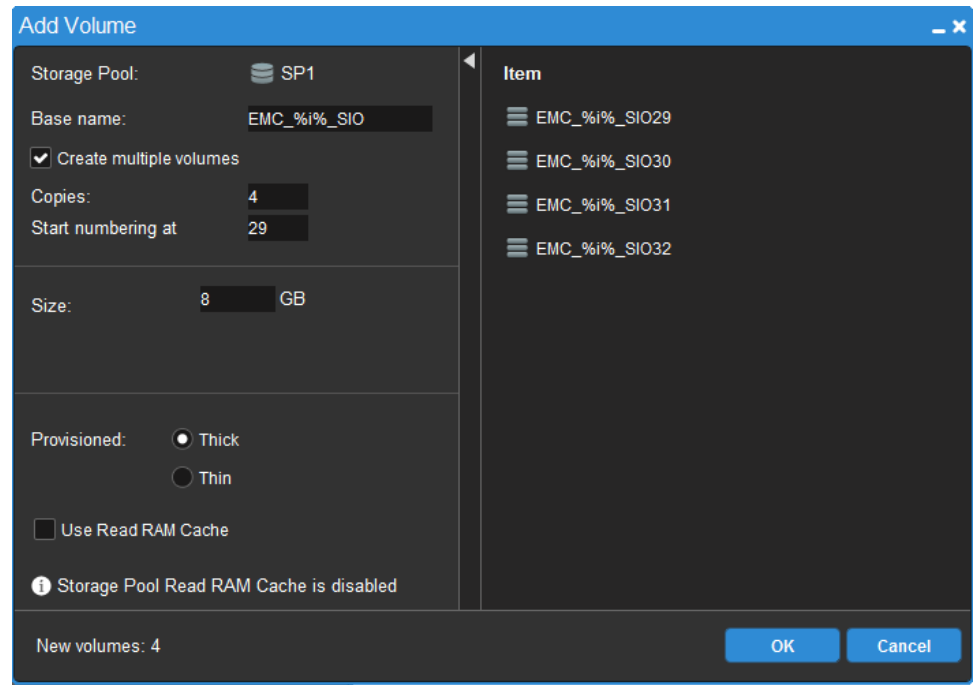
1. In any of the **Frontend > Volumes** views, navigate to the Storage Pool to which you want to add the volume, and select it.



2. From the **Command menu** or context-sensitive menu, select **Add Volume**.
3. In the **Add Volume** window, if you want to create more than one volume, select **Create multiple volumes** and type the number of volumes you would like to add in the **Copies** box.
  - If you type 1, only one volume will be created (optional—can be left blank).
  - 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 for the volume:
  - If you are adding one volume, enter the name in the **Name** box.
  - If you are adding multiple volumes, enter the base name in the **Base name** box.  
The volumes 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 you want to enable the RMcache feature (disabled by default), select **Use RMcache**.
9. 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 14 Add Volume window

**After you finish**

To use the created volume, you must map it to (at least) one SDC. If the restricted SDC mode is enabled for the system, you must approve SDCs prior to mapping volumes to them. For more information on approving SDCs, see [Approve SDCs \(GUI\)](#) on page 179. For more information on mapping volumes, see [Mapping a volume to an SDC](#) on page 90.

## 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 224. 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**

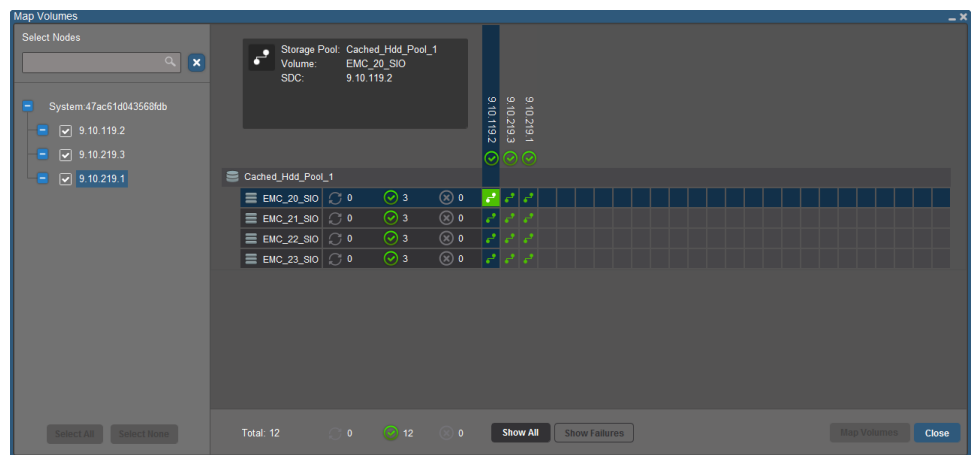
1. In the **Frontend > Volumes** view, navigate to the volumes, and select them.
2. From the **Command** menu or context-sensitive menu, select **Map Volumes**.

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 15** 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 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, "Managing System Objects"](#)

[Chapter 7, "Security Management"](#)

[Chapter 8, "Opening the GUI and Logging In"](#)

[Chapter 9, "GUI Features"](#)

[Chapter 10, "Monitoring the System using the GUI"](#)

[Chapter 11, "Configuring the System using the GUI"](#)

[Chapter 12, "Using the VMware Plug-in"](#)



# CHAPTER 6

## Managing System Objects

The following topics describe how to manage and configure ScaleIO system objects.

- [CLI basics](#)..... 96
- [Syntax](#)..... 97
- [Extend an existing ScaleIO system](#)..... 99
- [Managing the MDM cluster](#)..... 99
- [Managing the SDSs and cache](#)..... 109
- [SDC operations](#)..... 111
- [Managing ESX servers](#)..... 120

## 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
<code>--mdm_ip &lt;IP&gt;</code>	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 <code>--mdm_ip</code> parameter must be added to every CLI command.
<code>--approve_certificate</code>	Preemptive approval of the MDM's certificate
<code>&lt;command&gt;</code>	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 disable secure communications on every MDM server to enable execution of commands.

- To disable secure communications in Windows, on each MDM open the SCLI `conf.txt` file, and add the following line:

```
cli_use_secure_communication=0
```

- To disable secure communications in Linux, run the following on each MDM:

```
echo cli_use_secure_communication=0 >> ~/.scli/conf.txt
```

---

**Note**

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

---

## 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
--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	

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

### Example

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

## 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 31.

Proceed to the section that describes your environment:

- [Replacing a cluster member by adding a new member to the cluster](#) on page 100
- [Replacing a cluster member without adding a new server to the cluster](#) on page 102

## 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 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: 0x6dfelc5f4062b5b3
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. 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: 0x6dfelc5f4062b5b3
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: 0x6dfelc5f4062b5b3
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: 0x6dfelc5f4062b5b3
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 be 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: 0x6dfelc5f4062b5b3
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: 0x6dfelc5f4062b5b3
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\sdsc\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\sdsc\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 <i>ScaleIO User Guide</i> .
CLI	Add, modify, and remove virtual IP addresses.	For details, see the <i>ScaleIO CLI Reference Guide</i> .
REST API	Add, modify, and remove virtual IP addresses.	For details, see <i>ScaleIO REST API Reference Guide</i> .

## 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	<code>set_restricted_sdc_mode command</code>
Add an SDC to the approved list, when restricted SDC mode is enabled	<code>--add_sdc</code>

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 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 an 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: <code>/opt/emc/scaleio/sds/bin/close_firewall_port.sh</code>
Windows	From command line, run the batch file <code>C:\Program Files\EMC\scaleio\sds\bin\close_firewall_port.bat</code>

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	<code>/opt/emc/scaleio/sds/bin/port</code>
Windows	<code>C:\Program Files\EMC\scaleio\sds\bin\port</code>

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

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

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:  <pre>Pkill sds</pre>
Windows	From command line, run the command:  <pre>net stop sds_service &amp; net start sds_service</pre>

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, SSDs and NVMe 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

## Managing read RAM cache

This section describes how to manage the read RAM 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.

Read RAM 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 26.

---

### 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.

## 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 237.

## 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
<code>--ip &lt;MDM_IP_ADDRESS_LIST&gt;</code>	List of IP addresses (comma delimited) for this Master or Slave MDM
Optional:	
<code>--file &lt;CONFIG_FILE_NAME&gt;</code>	Name of the configuration file to which the MDM information should be written
<code>--only_cfg</code>	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 102 .

**Parameters**

Parameter	Description
<code>--ip &lt;EXISTING_MDM_IP_ADDRESS&gt;</code>	One of the existing MDM IP addresses
<code>--new_mdm_ip &lt;NEW_MDM_IP_ADDRESSES&gt;</code>	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:	
<code>--file &lt;CONFIG_FILE_NAME&gt;</code>	The name of the configuration file to which the MDM information should be written
<code>--only_cfg</code>	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 7

## Security Management

The following topics describe security management in ScaleIO.

- [Setting up SSH authentication on the ScaleIO Gateway](#)..... 124
- [Configuring SSL component authentication](#)..... 124
- [Managing SDC access to the MDM](#)..... 128
- [Approved encryption methods](#)..... 129
- [Login banner overview](#)..... 129

## 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 **System Settings** 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 126.

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

- **Linux:**  
/opt/emc/scaleio/gui/certificates
- **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	<code>set_restricted_sdc_mode</code> command
Add an SDC to the approved list, when restricted SDC mode is enabled	<code>--add_sdc</code>

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
<b>Create (or modify) a new banner</b>	<p>a. Create a text file (or modify an existing file) with the message that you want to display in the login banner.</p> <p>b. Run the following command:</p> <pre>scli --set_login_banner --filename &lt;FILENAME&gt;</pre> <p>where <i>&lt;FILENAME&gt;</i> is the path of the login banner text file.</p> <p>The login banner is displayed the next time a user logs into the ScaleIO system.</p>
<b>Stop displaying the banner</b>	<p>a. Run the following command:</p> <pre>scli --set_login_banner --remove_banner</pre>

## 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 8

## Opening the GUI and Logging In

The following topics describe how to open the GUI, and log in procedures

- [Log in to the ScaleIO GUI](#)..... 134
- [Connection and disconnection information](#)..... 134

## 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 "Install 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:
  - Linux: Run the script `/opt/emc/scaleio/gui/run.sh`.
  - Windows: Click **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 9

## GUI Features

The following topics describe GUI features.

- [GUI overview](#)..... 136
- [GUI conventions](#)..... 136
- [Dashboard view](#)..... 138
- [Frontend views](#)..... 146
- [Backend view](#)..... 149
- [Alerts view](#)..... 155
- [Property Sheets](#)..... 156

## GUI overview

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 many 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. Your user privileges control the features displayed by the GUI. If you cannot access certain commands, check whether you have the appropriate user permissions, and if necessary, contact your system administrator for assistance.

The following table provides a general overview of the tasks that you can perform with the GUI.

**Table 11** GUI task overview

To do this...	Use...
See a general overview of the entire system	Dashboard view—see <a href="#">Dashboard view</a> on page 138
See detailed information about one or more backend system objects, in table format (filtering available)	Backend view—see <a href="#">Backend view</a> on page 149
See detailed information about one or more frontend system objects, in table format (filtering available)	Frontend views—see <a href="#">Frontend views</a> on page 146
See very detailed information about a specific system object	Property Sheets—see <a href="#">Property Sheets</a> on page 156
Minimize the main window to a floating widget	Widget—see <a href="#">Widget</a> on page 140
See a list of errors and alerts currently active in the system	Alerts view—see <a href="#">Alerts view</a> on page 155
Add, remove, maintain, or configure a system object (backend)	Backend view—see <a href="#">Configuring the System using the GUI</a> on page 163
Add, remove, or configure volumes, snapshots or SDCs	Frontend views—see <a href="#">Configuring the System using the GUI</a> on page 163
Monitor various aspects of your system	Backend and Frontend views—see <a href="#">Monitoring the System using the GUI</a> on page 161
Find information about your system's license	System Settings menu—see <a href="#">Viewing licensing information</a> on page 162

## 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 155 and [ScaleIO Alerts in SNMP, GUI, REST, and ESRS](#) on page 277.

**Figure 16** 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 12** 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 12** 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 13** 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 17 Dashboard controls



- |   |  |
|---|--|
| 1 <b>Navigation toggle</b> —Displays navigation tree  | 6 <b>System Settings</b> menu—System Settings, Add Nodes, Upgrade, Log Collection, User Preferences, About   |
| 2 <b>Current view</b> —Displays object currently being applied to the dashboard   | 7 <b>Color code legend</b> —Displays the legend  |
| 3 <b>Dashboard</b> —Displays the Dashboard view, currently selected   | 8 <b>I/O Workload toggles</b> —control the information displayed in this tile. For more information, see <a href="#">Customizing system preferences</a> on page 201. |
| 4 Other tabs—Display various views in the GUI; some tabs provide further options when clicked                           | 9 <b>Dimmer toggle</b> —Enables\disables the dimmer  |
| 5 <b>Username</b> —name of user currently logged in to the GUI. To the right of the username is the <b>Logout</b> menu. | 10 <b>Widget toggle</b> —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.

## System Settings and Logout menus


The **Logout** menu lets you do the following:

- View the user name of the user currently logged in to the system (default)
- Log out of the system


The **System Settings** menu lets you do the following:

- 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 201.
- Display information about your system, including information required for licensing (**About** option). For more information, see [Viewing licensing information](#) on page 162.

## 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 18** 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 201.

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.

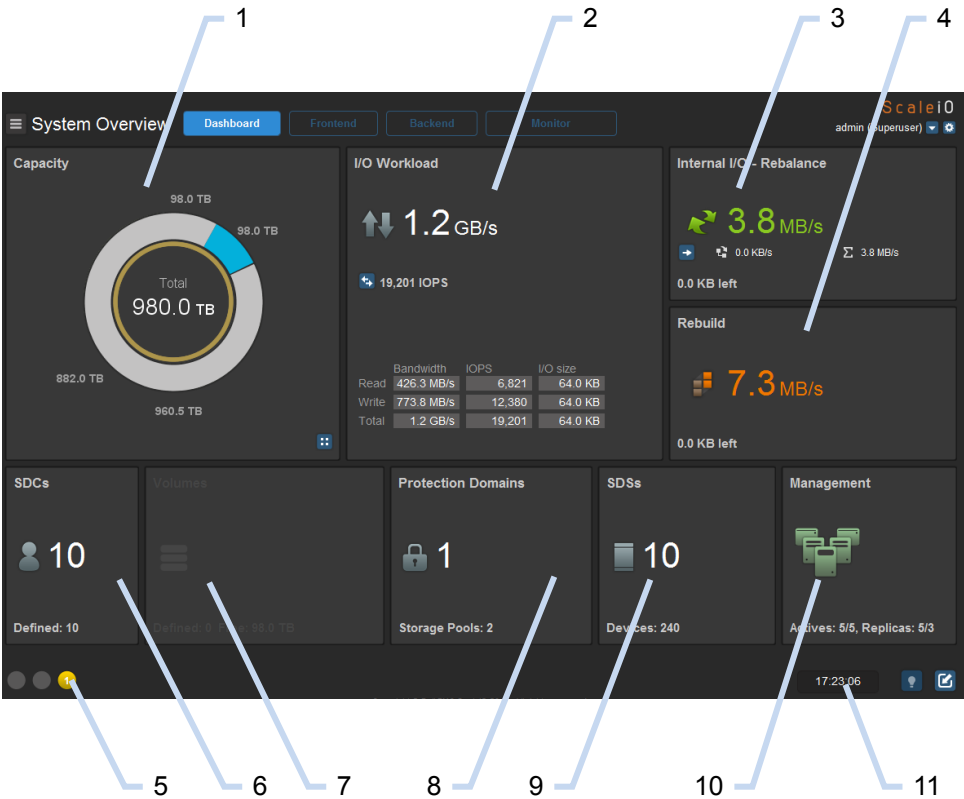
---

### Note

The filter used in the Backend view does not influence the Dashboard display.

---

Figure 19 Dashboard tiles



The following table describes the tiles displayed on the Dashboard.

Table 14 Dashboard tiles

Tile	Description
1 - Capacity	<p>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:</p> <p><b>Outer Ring</b></p> <p>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:</p> <ul style="list-style-type: none"><li><i>In use</i>—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.</li><li><i>Unused</i>—Shows how much further the raw capacity can be expanded; if this capacity is not accessible, it will be marked as Unavailable Unused.</li><li><i>Special purpose</i>—Shows spare capacity reserved for system operation; Decreased capacity that was deducted from devices (using the <code>Set Device Capacity Limit</code> command), and cannot be used.</li></ul> <p><b>Note</b></p> <p>A miniature version of the outer Capacity ring is shown in some Backend table views and Property Sheets.</p>

**Table 14** Dashboard tiles (continued)

Tile	Description
	<p><b>Center circle</b></p> <p>Displays the total amount of available raw storage.</p> <hr/> <p><b>Note</b></p> <p>Total available raw storage does not represent the total amount of capacity available for volume allocation.</p> <hr/> <p><b>Inner Ring</b></p> <p>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.</p>
2 - I/O Workload	<p>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: <b>Application I/O</b>, <b>Overall I/O</b>, and <b>I/O Bandwidth</b>.</p> <p>In Advanced Dashboard view (controlled by system Preferences—see <a href="#">Customizing system preferences</a> on page 201) aggregated values of bandwidth and IOPS are displayed.</p> <p>The table in this tile summarizes the Reads, Writes and Totals of IOPS, and throughput and the average size of an I/O.</p>
3 - Internal I/O	<p>The tile displays the system's internal I/O. Clicking on the arrow icon cycles through total internal I/O, internal I/O due to migration, and internal I/O due to rebalance operations.</p>
4 - Rebuild	<p>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.</p> <p>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 <a href="#">Customizing system preferences</a> on page 201) more details are displayed on this tile. Hovering over the title displays a tooltip listing the events that triggered the rebuild operation. Click <b>More details...</b> to open the <b>Monitor</b> view.</p>
5 - Alert Indicators	<p>Displays the number of active alerts in the system, using the system-wide color codes.</p>
6 - SDCs	<p>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).</p>
7 - Volumes	<p>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.</p> <p>Note that the number of volumes and the total capacity include snapshots.</p>
8 - Protection Domains	<p>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.</p>

**Table 14** 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 <a href="#">Management (MDM) cluster status</a> on page 144. 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 15** Management node icons with normal operational status (green)

Icon	Description
	5-node cluster
	3-node cluster
	single node

**Table 16** 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

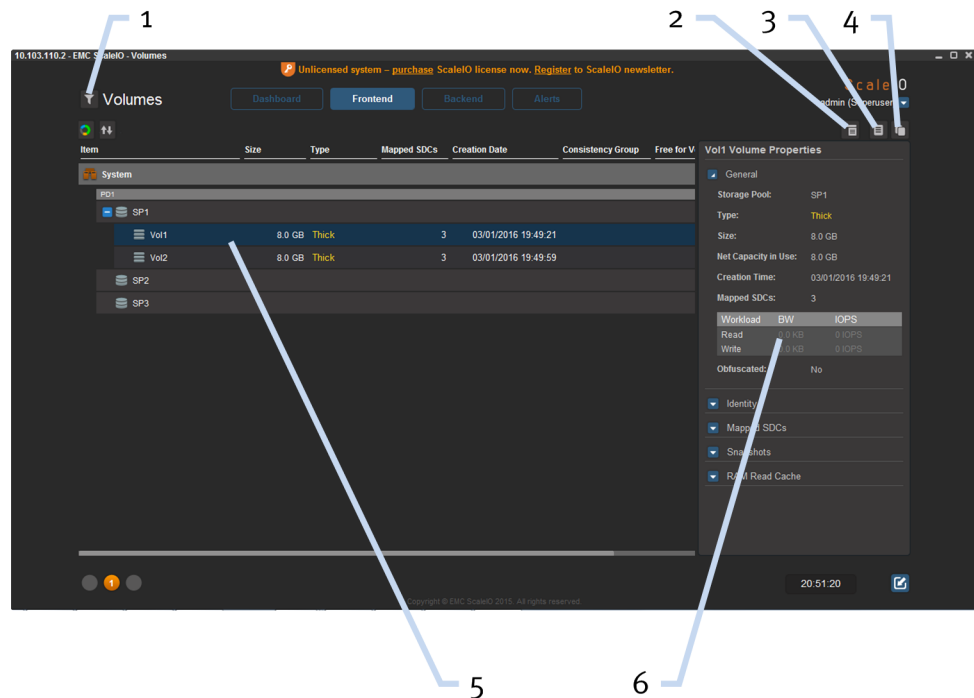
The **Frontend** view provides detailed information about frontend objects in the system, including volumes, SDCs and snapshots, and lets you perform various configuration operations. The main areas of the Frontend view are:

- **Filter**—lets you filter the information displayed in the table and Property Sheets. For more information, see [Filter \(1\)](#) on page 151.

- **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 156. For more information about the terminology used in the Property Sheets, see the Glossary at the end of this publication.

## Volumes

Figure 20 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 21 Frontend > SDCs view

Item	SDC IP Address	Connected	# Mapped Volumes	Read Bandwidth	Read IOPS	Read IO Size	Write Bandwidth	Write IOPS	Write IO Size	Alerts
System			2	0.0 KB/s	0	0.0 KB	0.0 KB/s	0	0.0 KB	
10.76.30.12	10.76.30.12	✓	2	0.0 KB/s	0	0.0 KB	0.0 KB/s	0	0.0 KB	
10.76.30.13	10.76.30.13	✓	2	0.0 KB/s	0	0.0 KB	0.0 KB/s	0	0.0 KB	
10.76.30.14	10.76.30.14	✓	2	0.0 KB/s	0	0.0 KB	0.0 KB/s	0	0.0 KB	

Snapshots


Figure 22 Frontend > filtered Snapshots view

Item (filtered)	Mapped
System	
Default_Protection_Domain	
Hdd_Pool_1	
rdm_vol_0	Yes
rdm_vol_1	Yes
rdm_vol_2	Yes
rdm_vol_3	Yes

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 151.
- **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 156. 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,  $\text{Max} = [\text{Total} + \text{decreased}]$  in the Backend table). Similarly, arrows and icons are unavailable in the same circumstances.

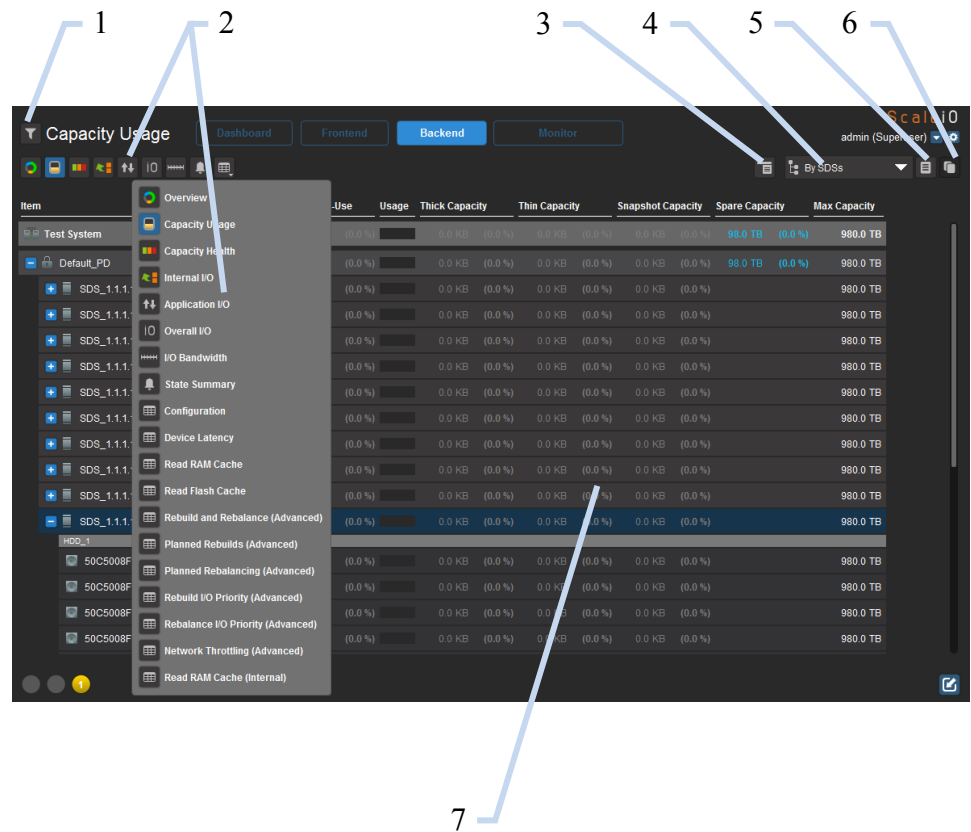
---

#### 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 23 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 151.
- 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 151.
- 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
- 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

sorting by SDSs, and Rebuild I/O  
Priority is only displayed when sorting  
by Storage Pools.

Command menu (3)

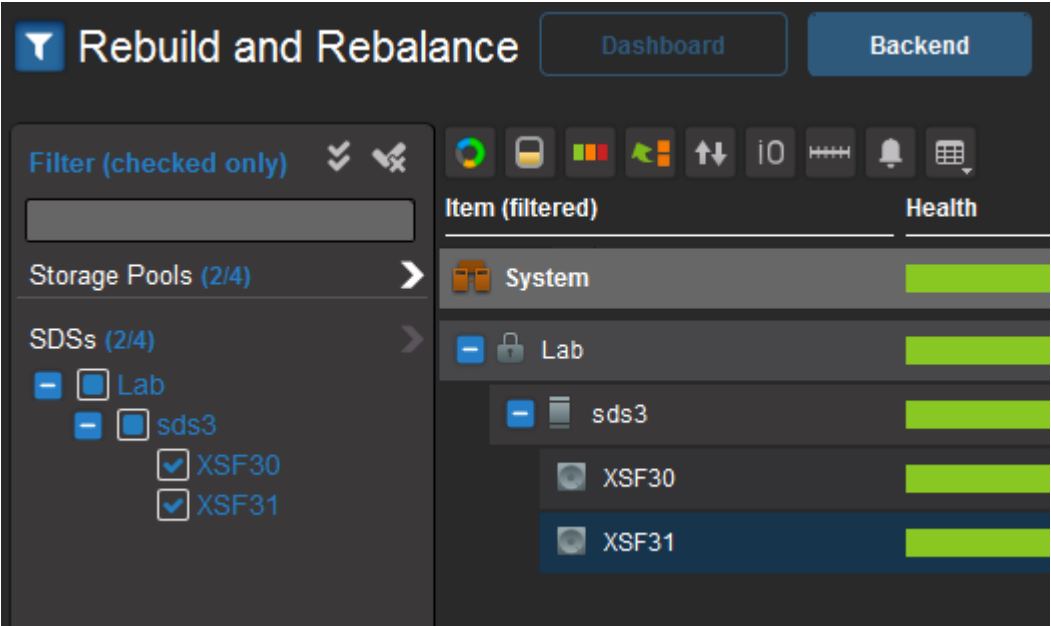
The **Command menu** button displays a list of commands which you can perform on rows selected in the table. The contents of the **Command menu** differ, depending on the object selected in the table. Many of the commands can also be accessed from the context-sensitive menu when table rows are right-clicked.

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 24 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 137 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 17** 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	<p>Provides a general overview the capacity and health of system objects.</p> <p>If you find that there are active alerts, you can switch to <b>State Summary</b> table view, or the <b>Alerts</b> view to see more details about the alerts.</p> <p>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.</p> <hr/> <p><b>Note</b></p> <p>The miniature Capacity ring shown here represents the outer Capacity ring on the Dashboard. For more information, see <a href="#">Dashboard tiles</a> on page 141.</p>
Capacity Usage 	Total, In-Use, Usage, Thick Thin, Snapshot, Spare, Max Related Property Sheet section: <b>Capacity</b>	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: <b>Capacity, Alerts, Rebuild/Rebalance</b>	Provides information about the health of the objects in the system, per object
Internal I/O 	Health, Backward Rebuild, Forward Rebuild, Rebalance Related Property Sheet sections: <b>Alerts, Rebuild/Rebalance</b>	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: <b>Workload</b>	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.





**Table 17** Table view options (continued)

Name of view	Contents of view	Suggested use, comments
Overall I/O 	Bandwidth, IOPS and I/O Size for: Total, Total Read, Total Write Related Property Sheet section: <b>Workload</b>	Provides workload information for all I/Os in the system, including both application I/Os, and I/Os for internal processes
I/O Bandwidth 	Read, Write, Backward Rebuild, Forward Rebuild, Rebalance, Total, Total Read, Total Write Related Property Sheet section: <b>Network Throttling</b>	Shows the bandwidth being used for various jobs in the system
State Summary 	Summary Related Property Sheet section: <b>Alerts</b>	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: <b>Identity, Related Objects</b>	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 <b>Free Capacity for Volume Allocation</b> 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 <a href="#">Configuring the System using the GUI</a> on page 163.  The <b>Related Objects</b> section of the Property Sheet helps you to identify related objects to the one selected in the table.
Device Latency 	Average Read Latency, Average Write latency, Average Read Size, Average Write Size, Scanned Capacity, Resolved Errors, Data Conflicts Related Property Sheet sections: <b>Device Latency , Background Device Scanner</b>	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.  <b>Note</b>  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.
Read RAM Cache 	State, Size, Used, Hit Rate, Write Mode Related Property Sheet section: <b>Read RAM Cache</b>	Lets you check which SDSs have Read RAM Cache enabled, view associated counters, and check which Storage Pool is set to use the cache

**Table 17** Table view options (continued)

Name of view	Contents of view	Suggested use, comments
		for its devices. Advanced feature; modify with caution.
Read Flash Cache 	State, Size, Used, Read Hit Rate, Total Errors, Alerts Related Property Sheet sections: <b>Read Flash Cache</b> and <b>Related Objects</b>	Lets you check which SDSs and Storage Pools have Read Flash Cache enabled, view associated counters. The <b>Related Objects</b> section of the corresponding Property Sheet shows the names/paths of the devices used for caching.
Rebuild and Rebalance (Advanced) 	Health, Backward Rebuild, Forward Rebuild, Rebalance Related Property Sheet section: <b>Rebuild/Rebalance</b>	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: <b>Rebuild/Rebalance</b>	<p>Shows amount of unprotected data, a visual breakdown of rebuild progress (Distribution), bandwidth, direction and status (Active/Pending) of jobs.</p> <p>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.</p> <hr/> <p><b>Note</b></p> <p>Overall Degraded capacity includes both healthy and failed copies of the data, while the processed capacity includes only the failed copies.</p> <hr/> <p>Advanced feature; modify with caution.</p>
Planned Rebalancing (Advanced) 	Protected, Distribution, Rebalance Related Property Sheet section: <b>Rebuild/Rebalance</b>	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: <b>I/O Priority</b>	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: <b>I/O Priority</b>	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,	Displays Network Throttling settings currently configured in the system. These settings control

**Table 17** Table view options (continued)

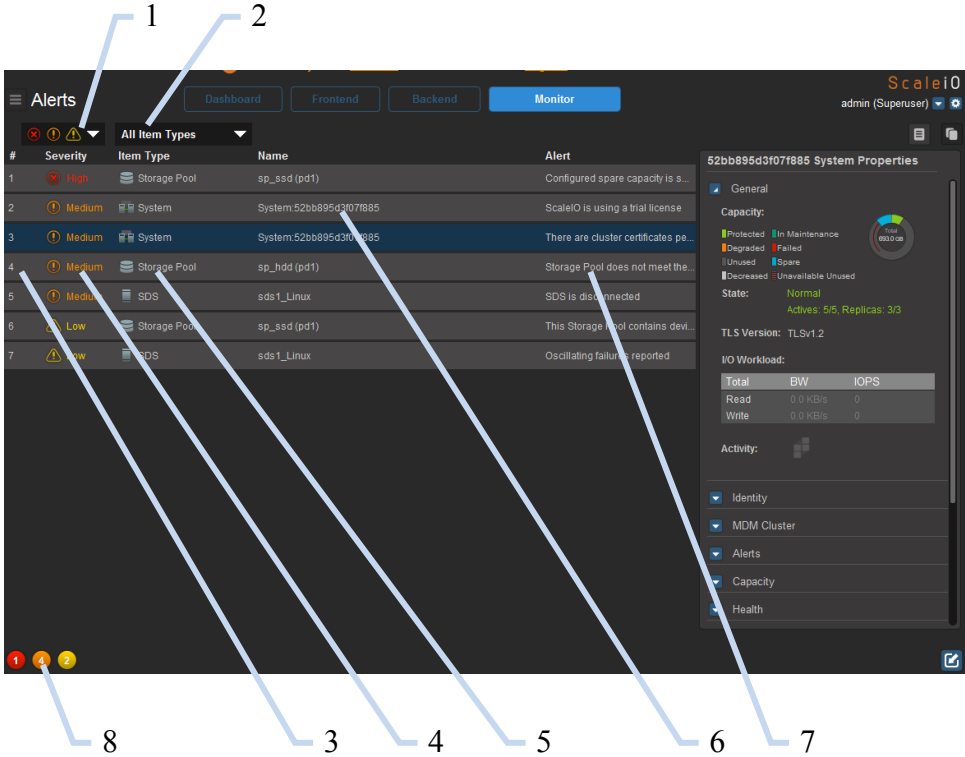
Name of view	Contents of view	Suggested use, comments
	Rebalance Queue Length, Rebuild Outgoing Jobs, Rebalance Outgoing Jobs Related Property Sheet section: <b>Network Throttling</b>	system performance. Advanced feature; modify with caution.
Read RAM Cache (Internal) 		For support purposes only. Visible only if Backend Internals are enabled in the <b>User Preferences</b> 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 277

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 156.

**Figure 25** Alerts view



- 1 Severity filter—filters the table contents according to alert severity:
- 5 **Item Type**—MDM, Protection Domain, Storage Pool, SDS, Device, SDC

**All alerts**

**Medium**—alerts which are Medium or High severity

**High**—only alerts which are High severity

- |  |  |
|--|--|
| <p>2 <b>Item Types filter</b>—filters the table contents according to:</p> <p>All Item Types, MDM, Protection Domain, Storage Pool, SDS, Device, SDC</p> | <p>6 <b>Name</b>—the user-defined name of the item, if one has been defined</p>  |
| <p>3 <b>#</b>—the line number in the Alerts table</p>  | <p>7 <b>Alert</b>—the alert message generated by the system</p>  |
| <p>4 <b>Severity</b>—the alert severity: High, Medium or Low</p>   | <p>8 <b>Alert indicators</b>—summarize the total amount of each alert type (from left to right: High, Medium, Low)</p> |
- 

## Property Sheets

The Property Sheets provide detailed read-only information about the object selected in most **Frontend** tables, **Backend** tables, **Monitor** tables.

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 Volumes** for the selected SDC, including the volume name, and bandwidth limit and IOPS limit per volume
- **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.

---

- **Read RAM 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.
- **Certificate Info** is shown for security certificates
- **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.

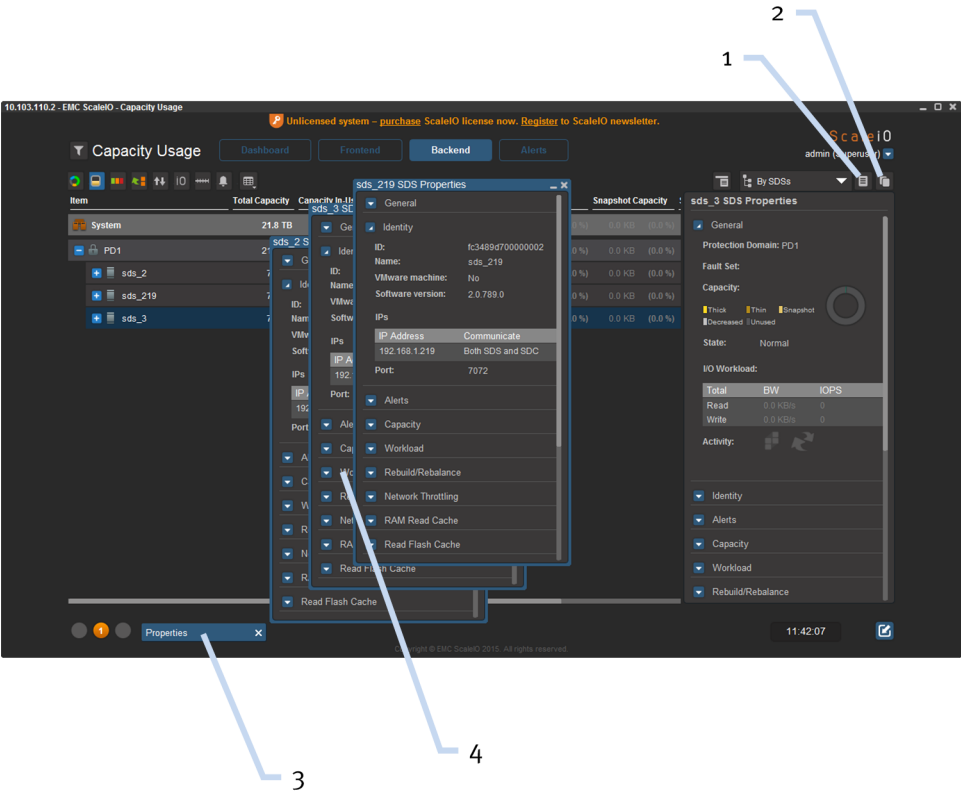
---

**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 26 Multiple floating Property Sheets



- 1

Display\Hide Property Sheet button
- 2

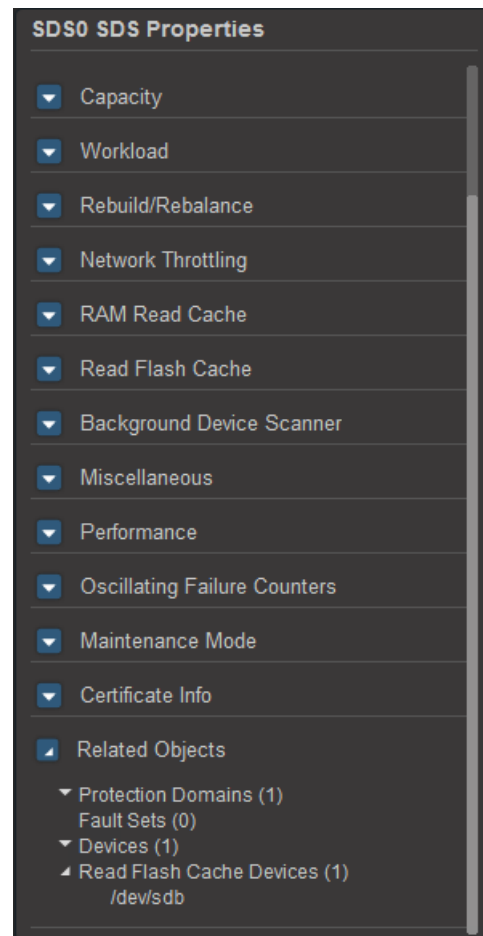
Duplicate Property Sheet button
- 3

Duplicated Property Sheets toggle button
- 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 27** Example of a Property Sheet for an SDS





# CHAPTER 10

## Monitoring the System using the GUI

The following topics explain how to monitor the system, using the GUI.

- [Viewing object properties in Backend and Frontend views](#).....162
- [Viewing licensing information](#).....162
- [Verifying your connection to the Management cluster](#)..... 162

## 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 156.

---

## 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 **System Settings** 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 144.

If the connection is lost, the GUI is dimmed, and a dialog box is displayed.

# CHAPTER 11

## Configuring the System using the GUI

The following topics explain how to configure the system, using the GUI.

• <a href="#">Configuring capacity</a> .....	164
• <a href="#">Configuring cache</a> .....	171
• <a href="#">Configuring volumes, volume trees, SDCs, and snapshots</a> .....	176
• <a href="#">Entering and exiting SDS Instant Maintenance Mode</a> .....	188
• <a href="#">Configuring Oscillating Failure counters</a> .....	189
• <a href="#">Applying Performance Profiles to system components</a> .....	194
• <a href="#">Configuring I/O priorities and bandwidth use (advanced)</a> .....	195
• <a href="#">Enabling and disabling Rebuild/Rebalance (advanced)</a> .....	197
• <a href="#">Using the background device scanner</a> .....	197
• <a href="#">Modifying Checksum protection mode</a> .....	199
• <a href="#">Renaming objects</a> .....	200
• <a href="#">Approving pending security certificates</a> .....	201
• <a href="#">Customizing system preferences</a> .....	201

## 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 33.

---

#### 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 171 and [Setting Read Flash Cache policy at SDS level](#) on page 172.

---

#### 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.

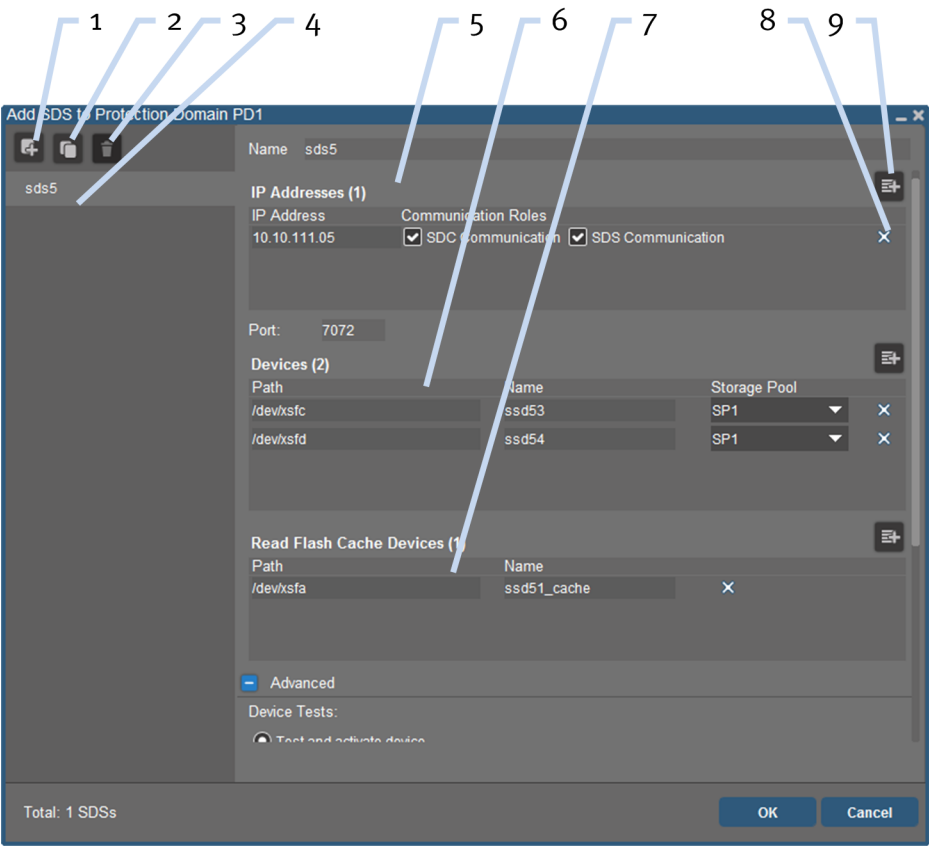
4. Click **OK**.

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.

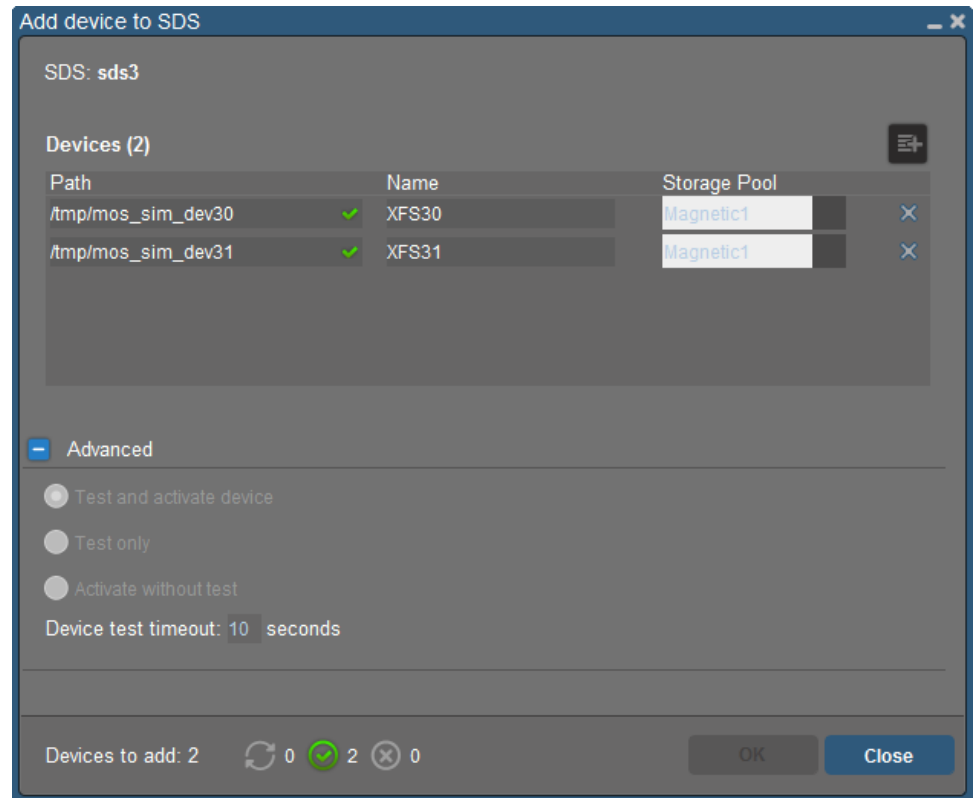
5. Click **Close**.

If you chose the **Test only** option in step 3, activate the devices as described in [Activating devices](#) on page 170.

Figure 28 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 29** 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 rebuild and rebalance operations in the system later on. For more information, see [Setting device capacity limits](#) on page 170.

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 > Storage** view, navigate to the desired object in the table, and select its row.
2. Right-click the row and 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. There 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 restarting 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.

---

### Procedure

1. To add a Protection Domain, perform these steps: In the **Backend > Storage** view, select the **System** row.
  - a. Right-click the row and select **Add Protection Domain**.  
The **Add Protection Domain** window is displayed.
  - b. 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, Fault Sets, Storage Pools, and Acceleration Pools to the Protection Domain. Before you add devices, ensure that at least one suitable Storage Pool is defined in the Protection Domain.
2. To remove Protection Domains, perform these steps:
  - a. In the **Backend > Storage** view, navigate to, and select one or more Protection Domains.



- b. Verify that you have removed all child nodes from the Protection Domain.
  - c. Right-click the Protection Domain and select **Remove**.
  - d. Click **OK**.
  - e. If a confirmation window appears, confirm the operation, and type your password if requested to do so.
3. To inactivate Protection Domains, perform these steps:
    - a. In the **Backend > Storage** view, navigate to, and select one or more Protection Domains.
    - b. Right-click the Protection Domain and select **Inactivate Protection Domain**.  
The **Inactivate Protection Domain** window is displayed.
    - c. Click **OK**.  
If a confirmation window appears, confirm the operation, and type your password if requested to do so.
  4. To activate Protection Domains, perform these steps:
    - a. In the **Backend > Storage** view, navigate to, and select one or more Protection Domains.
    - b. Right-click the Protection Domain and select **Activate Protection Domain**.  
The **Activate Protection Domain** window is displayed.
    - c. 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 35. Failure to do so may prevent creation of volumes.

---

### Procedure

1. In the **Backend > Storage** view, navigate to, and select the Protection Domain.
2. Right-click the Protection Domain and 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.

## Adding Storage Pools

A Storage Pool is a group of devices within a Protection Domain. Create Storage Pools before you start adding devices to the system. Each time that you add devices to the system, you must map them to either Storage Pools or Acceleration Pools.

---

**Note**

You cannot enable zero padding after adding the devices. For more information, see [Storage Pools](#) on page 33.

---

**Procedure**

1. In the **Backend > Storage** view, select the desired Protection Domain.
  2. Right-click the Protection Domain and select **Add Storage Pool**.  
The **Add Storage Pool** window is displayed.
  3. Type a name in the **Name** box.
  4. Select a media type.  
All devices added to this Storage Pool must be this type of device.
  5. Select an external acceleration type.
  6. If you want to use Read RAM Cache, select the corresponding checkbox and the desired Write Handling Mode.
- 

**Note**

The Read RAM 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 Read RAM Cache` command. For more information about Read RAM Cache features, see [Managing read RAM cache](#) on page 111.

---

7. Click **OK**.

## Activating devices

Use the `Activate Device` command in the following situations:

- Storage devices were added to the system using the **Test only** option for Device Tests, and successfully passed the tests.
- Storage devices were inactivated, and you want to bring them back online.

**Procedure**

1. In the **Backend > Storage** view, navigate to the device or devices in the table, and select the corresponding rows.
2. Right-click and select **Activate Device**.

## Clearing device errors

**Procedure**

1. In the **Backend > Storage** view, navigate to the device in the table, and select its row.
2. Right-click the row and select **Clear Device Errors**.

## Setting device capacity limits

In circumstances where you need to replace a storage device in your system with a storage 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 storage 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** > **Storageview**, navigate to the device in the table, and select its row.
2. Right-click the row and 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 171
- [Setting Read Flash Cache policy at SDS level](#) on page 172
- [Adding Read Flash Cache devices](#) on page 172
- [Removing Read Flash Cache devices](#) on page 174
- [Changing Read RAM Cache volume settings](#) on page 173
- [Configuring Read RAM Cache \(advanced, Backend\)](#) on page 174

### 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.

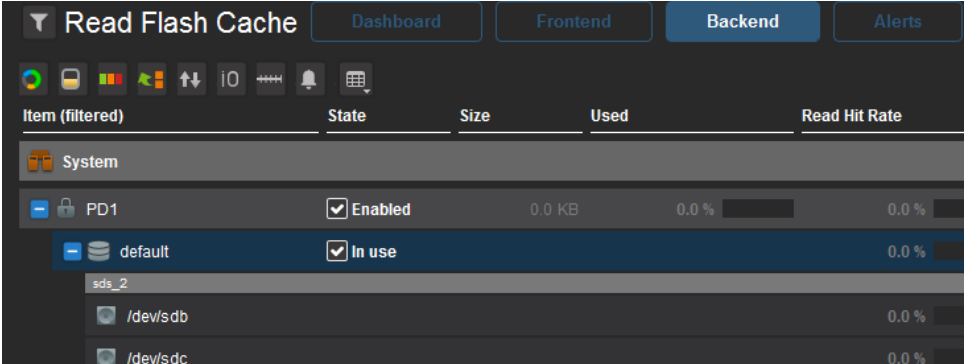
- Once this is enabled at Storage Pool level, set the Read Flash Cache policy at SDS level.
- Once Read Flash Cache is enabled at both Storage Pool and SDS levels, add Read Flash Cache devices to the SDS in order to commence caching.

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.



Item (filtered)	State	Size	Used	Read Hit Rate
System				
PD1	Enabled	0.0 KB	0.0 %	0.0 %
default	In use			0.0 %
sds_2				
/dev/sdb				0.0 %
/dev/sdc				0.0 %

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:

Item	State	Size	Used	Read Hit Rate	Total Errors	Alerts
System	Enabled	0.0 KB	0.0 %	0.0 %	0	1
PD1	Enabled	0.0 KB	0.0 %	0.0 %	0	0
sds_2	Enabled	0.0 KB	0.0 %	0.0 %	0	0
default	Enabled	0.0 KB	0.0 %	0.0 %	0	0
/dev/sdb	Enabled	0.0 KB	0.0 %	0.0 %	0	0

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 RCache 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 Read RAM Cache volume settings

By default, Read RAM Cache is disabled on volumes. To change Read RAM Cache settings on volumes, perform these steps:


**Procedure**

1. In **Frontend > Volumes**, select the **Volumes** or **Volumes Monitor** view.
2. Navigate to the volumes, and select them.
3. Right-click the volumes and select **Set Volume Read RAM Cache**.

The **Set Volume Use Read RAM Cache** window is displayed, showing a list of the volumes that will be modified.

4. Select or clear the **Use Read RAM Cache** check box as follows:
  - To disable Read RAM Cache on the volumes, clear the check box.
  - To enable Read RAM Cache on the volumes, select the check box.
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 30** Set Volume Use Read RAM 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 Read RAM 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 the *ScaleIO CLI Reference Guide*.

To configure caching, perform these steps:

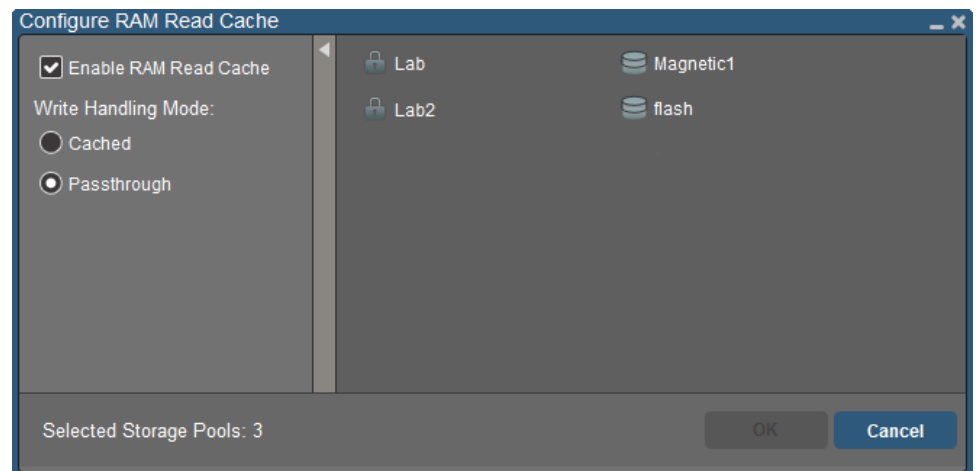
#### Procedure

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 31** 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 32** 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 explains how to remove volumes, create snapshots, and set volume bandwidth and IOPS limits. It also describes setting the SDC restriction mode and how to approve SDCs before mapping volumes.

### Adding volumes

Add volumes to a system.

#### Before you begin

There must be at least three SDS nodes in the system and there must be sufficient capacity available.

---

#### Note

For the minimum size of an SDS, see [System requirements](#) on page 20.

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 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 Read RAM Cache volume settings](#) on page 173.

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 any of the **Frontend > Volumes** views, 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**.
3. In the **Add Volume** window, if you want to create more than one volume, select **Create multiple volumes** and type the number of volumes you would like to add in the **Copies** box.
  - If you type 1, only one volume will be created (optional—can be left blank).
  - 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 for the volume:
  - If you are adding one volume, enter the name in the **Name** box.
  - If you are adding multiple volumes, enter the base name in the **Base name** box.  
The volumes 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 you want to enable the RMcache feature (disabled by default), select **Use RMcache**.
9. 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 33** Add Volume window

**After you finish**

To use the created volume, you must map it to (at least) one SDC. If the restricted SDC mode is enabled for the system, you must approve SDCs prior to mapping volumes to them. For more information on approving SDCs, see [Approve SDCs \(GUI\)](#) on page 179. For more information on mapping volumes, see [Mapping a volume to an SDC](#) on page 90.

**Restricted SDC mode**

Enabling restricted SDC mode gives you more control over your ScaleIO system.

Restricted SDC mode is set at the system level. When enabled, you must approve SDCs prior to mapping volumes to them. The restricted SDC setting has the following modes:

- No restriction — Volumes can be mapped without pre-approving SDCs.
- GUID restriction — SDCs are approved using their GUID. Only once the SDC is approved can it be used for mapping volumes.
- Approved IP restriction — SDCs must be approved using their GUID and IP address. Only once the SDC is approved using both GUID and IP address can it be used for mapping volumes.

You can set the restricted SDC mode using the GUI, CLI, or REST API.

**Note**

In a system that has been upgraded and already has volumes mapped to SDCs, if you want to enable restricted SDC mode, you must first approve the SDCs and only then enable restricted SDC mode.

## Set the system's restricted SDC mode (GUI)

Use the ScaleIO GUI to set the restricted SDC mode.

The system's restricted SDC mode can also be set using the CLI or the REST API. For details, see the *CLI Reference Guide* or the *REST API Reference Guide*.

---

### Note

In a system that has been upgraded and already has volumes mapped to SDCs, if you want to enable restricted SDC mode, you must first approve the SDCs and only then enable restricted SDC mode.

---

### Procedure

1. In the **Frontend > SDCs** view, right-click on the System and select **Set Restricted SDC Mode**.
2. In the **Set Restricted SDC Mode** dialog box, select one of the following options and click **OK**.
  - No restriction
  - GUID restriction
  - Approved IP restriction

### Results

The restricted SDC mode is set for the system. If you enabled restricted SDC mode by selecting either `GUID restriction` or `Approved IP restriction`, you must configure approved SDCs before you can map volumes.

## Approve SDCs (GUI)

When the system's restricted SDC mode is set to GUID restriction, you must approve SDCs before you can map volumes.

### Procedure

1. In the **Frontend > SDCs** view, right-click on one or several SDCs that you want to approve and select **Approve SDC**.
2. In the **Approve SDC** window, verify that the SDCs listed are the ones you want to approve and click **OK**.

### Results

The SDCs are approved and you can map volumes. The **Approved IPs** column in the **Frontend > SDCs** displays which SDCs are approved.

## Configure approved SDC IP addresses (GUI)

When the system's restricted SDC mode is set to `approved IP restriction`, you must configure SDC IP addresses before you can map volumes.

### Before you begin

Ensure that the SDCs have been approved by GUID.

### Procedure

1. In the **Frontend > SDCs** view, right-click on the SDCs that you want to approve and select **Configure Approved IP Addresses**.
2. In the **Configure Approved IP Addresses** window, add the IP addresses of the SDCs you want to approve.

You can click the **Add IP Address** button to add up to a total of four IP addresses.

3. Click **OK** and then click **Close**.

### Results

The SDC IP addresses are approved and you can map volumes. The **Approved IPs** column in the **Frontend > SDCs** view displays which SDC are approved.

## 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.

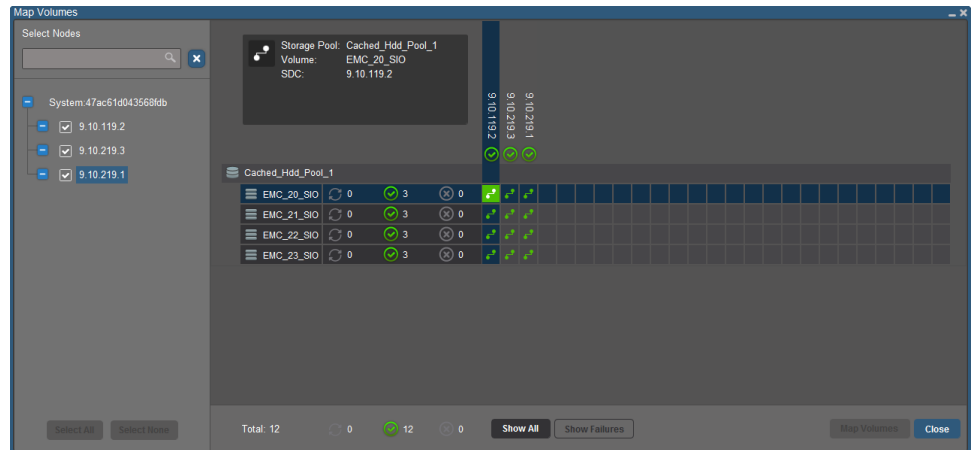
To identify the unique ID, run the `ls -l /dev/disk/by-id/` command. For more information, see [Associating ScaleIO volumes with physical disks](#) on page 224. You can also identify the unique ID using VMware. In the VMware management interface, device is called **EMC Fibre Channel Disk**, followed 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 34** 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**.

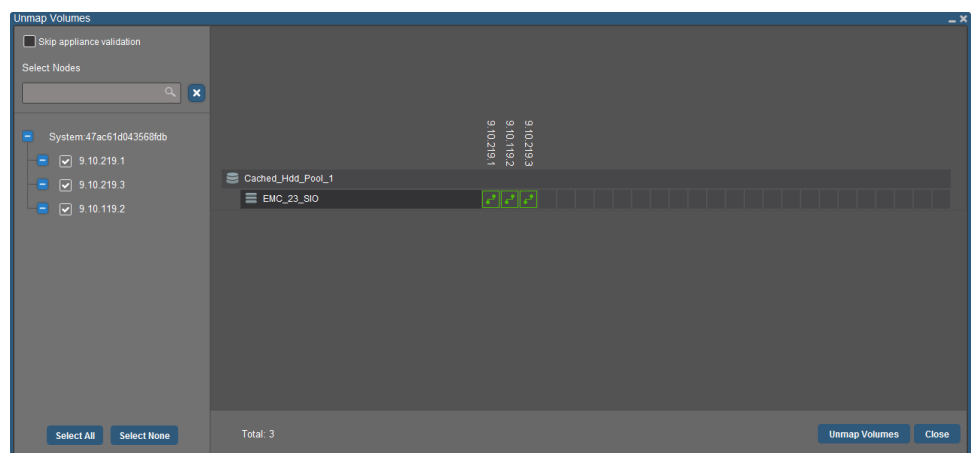
The **Unmap Volumes** window is displayed, showing a list of the volumes that will be unmapped.

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 35** 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 180.

If you want to remove a volume's related snapshots as well, or just the snapshots, see [Removing snapshots](#) on page 183. 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 36** 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

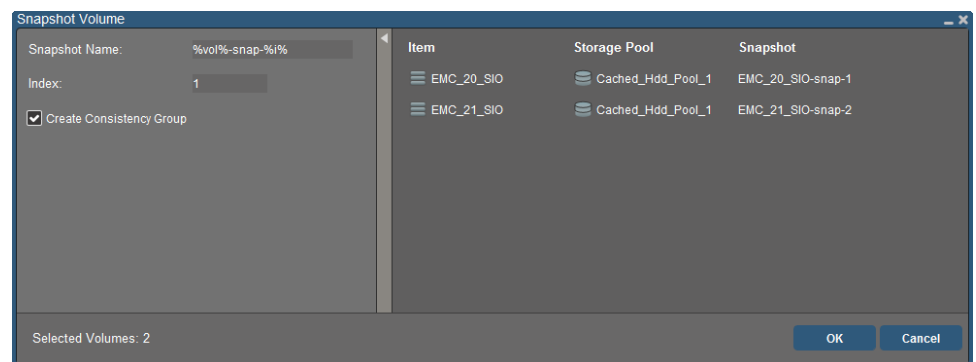
1. In the **Frontend > Volumes** view, navigate to the volumes, and select them.
2. From the **Command menu** or context-sensitive menu, select **Snapshot Volume**.

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 37** 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 180.

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 38** Removal of a volume and associated snapshots**Figure 39** 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 40** 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 41** 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:

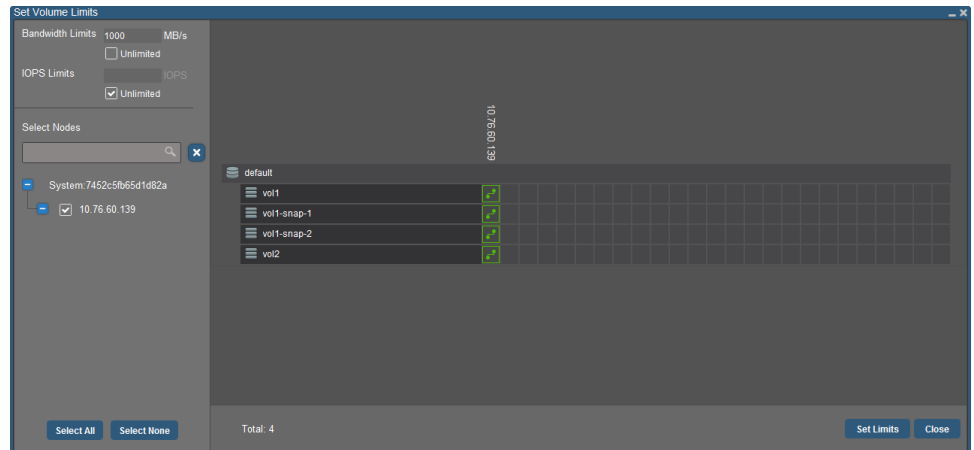
### Procedure

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 42** 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.

### Procedure

1. To put an SDS into Maintenance Mode, perform these steps:
  - a. In the **Backend > Storage** view, navigate to, and select the desired SDS.
  - b. From the **Command** menu or context-sensitive menu, select **Enter Maintenance Mode**.  
The **Enter Maintenance Mode** window is displayed.
  - c. 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

d. 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).

2. To put an SDS back into regular service (cancel Maintenance Mode), perform these steps:

- a. In the **Backend > Storage** view, navigate to, and select the desired SDS.
- b. From the **Command menu** or context-sensitive menu, select **Exit Maintenance Mode**.

The **Exit Maintenance Mode** window is displayed.

- c. If you want to force exit from Maintenance Mode even though there is a failed SDS, select the **Force Failed SDS** check box.

d. 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 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 <b>Backend &gt; Storage</b> view, select the <b>System</b> icon.
To configure counter parameters for a specific Protection Domain or Storage Pool:	In the <b>Backend &gt; Storage</b> 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:

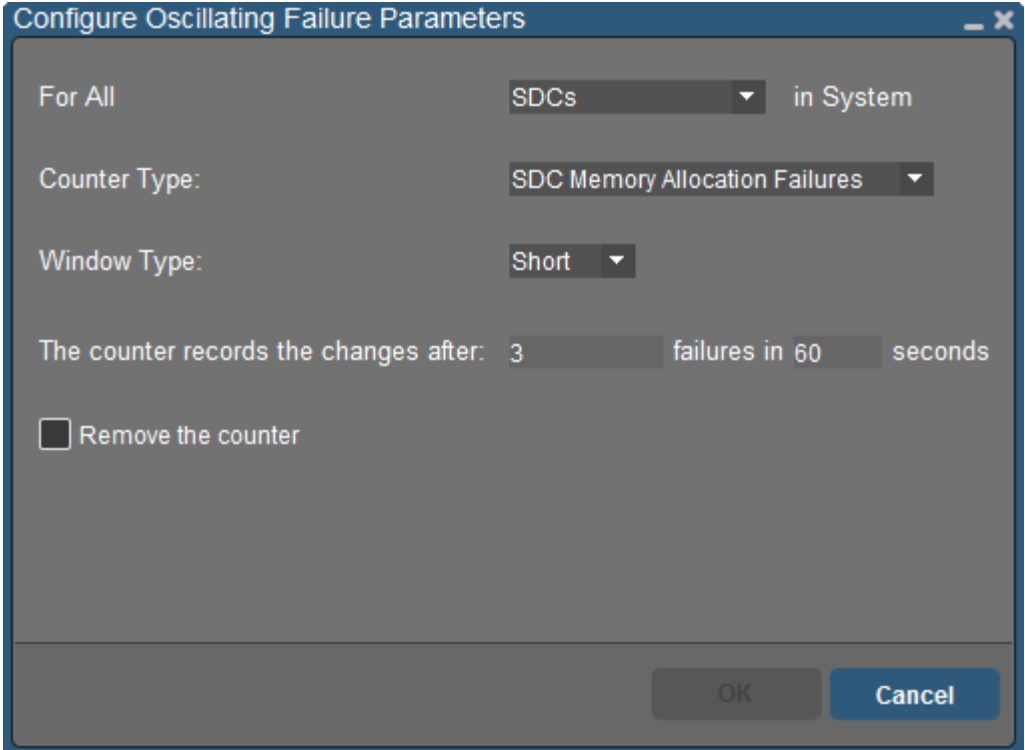
Option	Description
<b>For system level:</b>	In the <b>For All</b> box, select an option: <b>SDCs</b> , <b>Protection Domains</b> , or <b>Storage Pools</b> .
<b>For a Protection Domain or a Storage Pool:</b>	Go to the next step.

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
<b>If you want to remove the selected counter definition from the system:</b>	Select the <b>Remove the counter</b> check box.
<b>If you want to modify the threshold for the selected counter definition:</b>	Enter a number in the fields for: <ul style="list-style-type: none"> <li>• <b>failures</b> (the maximum number of failures per time interval before reporting begins)</li> <li>• <b>seconds</b> (the number of seconds per time interval)</li> </ul>

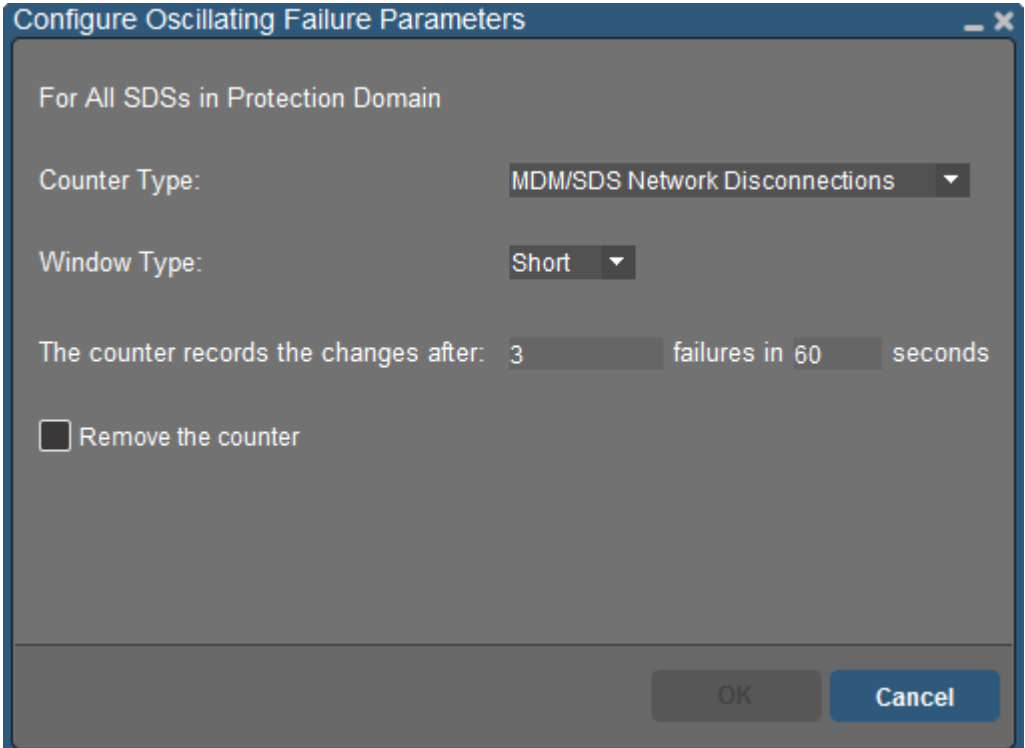
7. Click **OK**.

The currently configured counter parameters are displayed in the corresponding Property Sheet, in the **Oscillating Failure Parameters** section.

**Figure 43** Configure Oscillating Failure counters—System

The dialog box is titled "Configure Oscillating Failure Parameters". It contains the following fields and controls:

- For All**: A dropdown menu showing "SDCs" and the text "in System".
- Counter Type:**: A dropdown menu showing "SDC Memory Allocation Failures".
- Window Type:**: A dropdown menu showing "Short".
- The counter records the changes after:** A text input field with "3", followed by the text "failures in", another text input field with "60", and the text "seconds".
- Remove the counter:** A checkbox that is currently unchecked.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

**Figure 44** Configure Oscillating Failure counters—Protection Domain or Storage Pool

The dialog box is titled "Configure Oscillating Failure Parameters". It contains the following fields and controls:

- For All SDSs in Protection Domain**: A text label.
- Counter Type:**: A dropdown menu showing "MDM/SDS Network Disconnections".
- Window Type:**: A dropdown menu showing "Short".
- The counter records the changes after:** A text input field with "3", followed by the text "failures in", another text input field with "60", and the text "seconds".
- Remove the counter:** A checkbox that is currently unchecked.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

## 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, go to step 1.
- To reset counters for a specific Protection Domain or Storage Pool, go to step 6 on page 192.

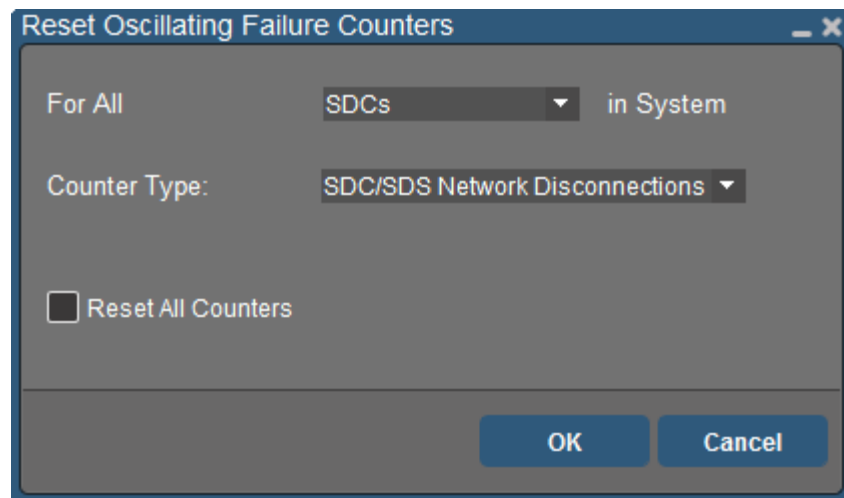
### Procedure

1. In the **Backend > Storage** view, select the **System** icon.
2. Right-click and 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. Perform one of the following:
  - For a specific counter, in the **Counter Type** box, select the required counter.
  - For all counters, in the **Counter Type** box, select **None**, and select the **Reset All Counters** check box.
5. Click **OK**.

**Figure 45** Reset Oscillating Failure counters—System



To reset counters for a specific Protection Domain or Storage Pool, perform the following steps:

6. In the **Backend > Storage** view, navigate to, and select the desired Protection Domain or Storage Pool.
7. Right-click and select **Reset Oscillating Failure Counters**.
8. Perform one of the following:
  - For a specific counter, in the **Counter Type** box, select the required counter.



- For all counters, select the **Reset All Counters** check box.
9. Click **OK**.

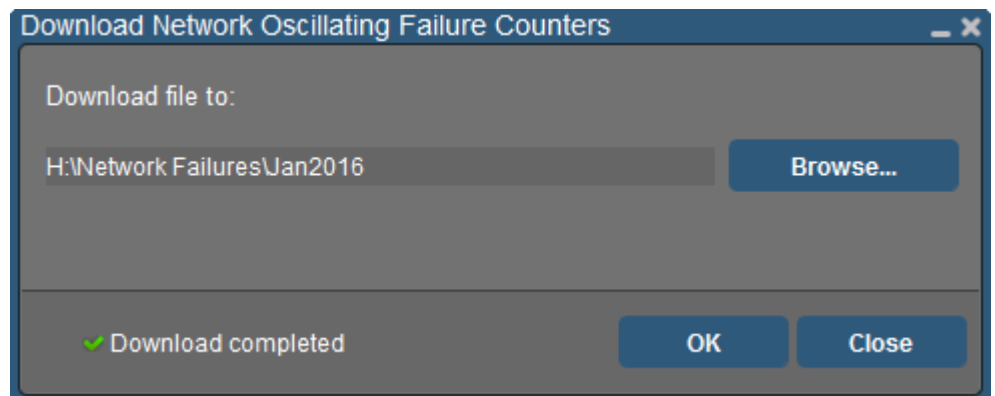
**Figure 46** Configure Oscillating Failure counters—Protection Domain or Storage Pool



## 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 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. Depending on the system component that you want to configure, in either the **Backend > Storage** or **Frontend > SDCs** view, navigate to, and select the desired objects.

**Note**

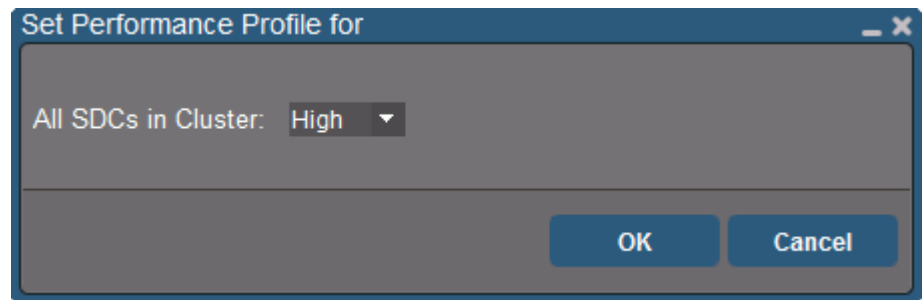
If you want to apply the Performance Profile to MDMs, select the System object.

2. Right-click the object and 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 47** 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 > Storageview**, navigate to, and select the desired Storage Pool.
2. Right-click the Storage Pool and 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 > Storage** view, navigate to, and select the desired Storage Pool.
2. Right-click the Storage Pool and select **Set I/O Priority**.
3. Select the desired options and edit values, and click **OK**.

## Configuring Network Throttling

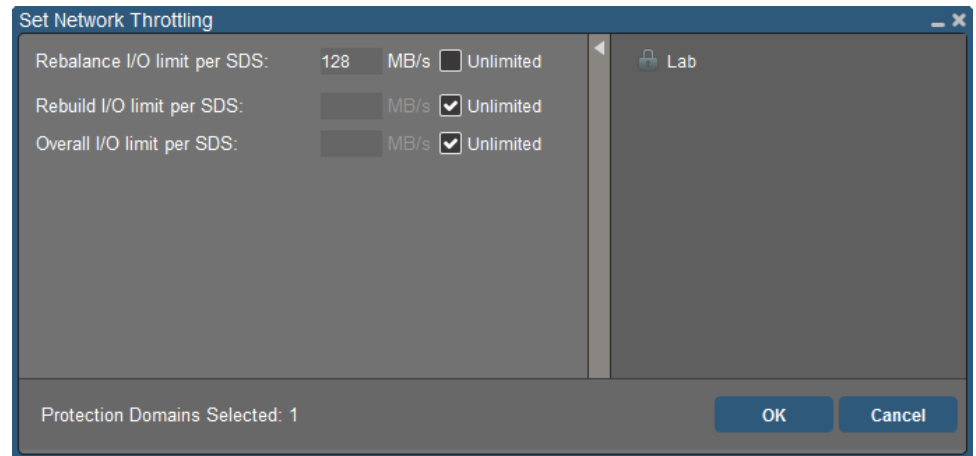
Network throttling affects network limits, and is used to control the flow of traffic over the network. It is configured per Protection Domain. The SDS nodes transfer data between themselves. This data consists of user-data being replicated as part of the RAID protection, and data copied for internal 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 > Storage** view, navigate to the desired Protection Domain, and select its row in the table.
2. Right-click the row and select **Set Network Throttling**.  
The **Set Network Throttling** window is displayed.
3. Configure the settings, and click **OK**.

**Figure 48** 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 > Storage** view, navigate to, and select the desired Storage Pools.
2. Right-click the Storage Pool and 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 251.

## 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 > Storage** 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, perform one of the following:
  - To enable the scanner, select the check box, and proceed to the next step.
  - 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).

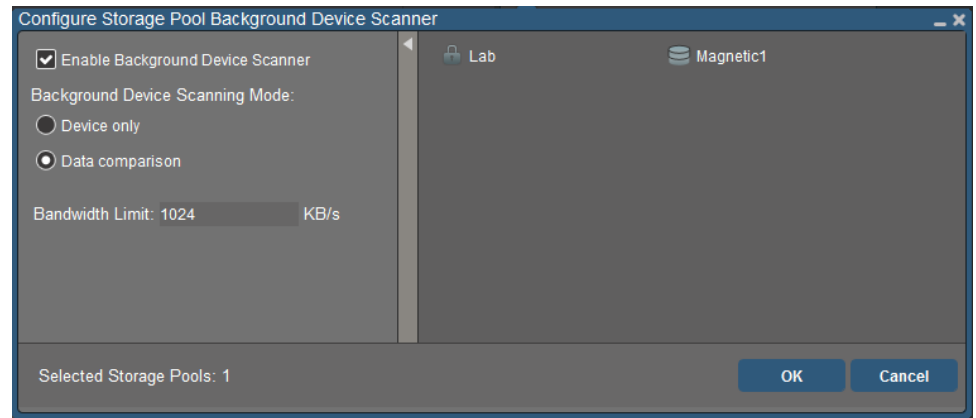
---

### 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 49** 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.

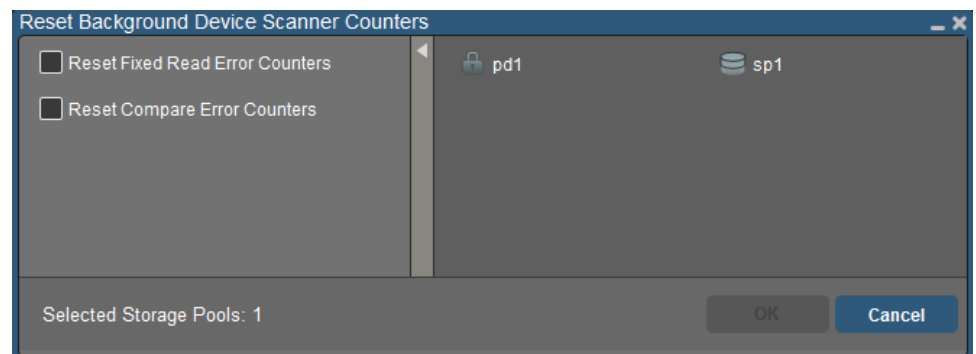
To reset counters, follow these steps:

### Procedure

1. In the **Backend > Storage** view, navigate to, and select the desired Storage Pools.
2. Right-click the Storage Pools and 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 50** 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:

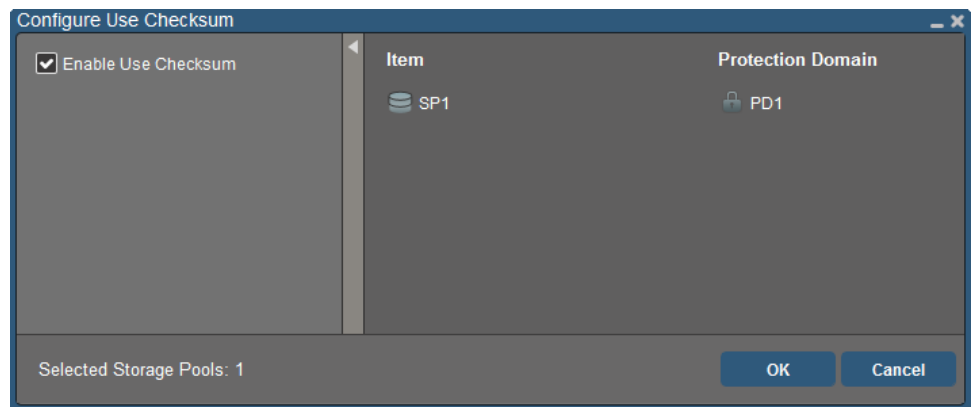
### Procedure

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 51** 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 > Storage** or any of the **Frontend** views, navigate to the object in the table, and select its row.
2. Right-click the object and select **Rename**.  
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

Approve pending security certificates, and view approved certificates in the **System Settings** window.

---

**Note**

When there are pending certificates in the system, they are listed in the **Backend > Storage** view > **State Summary** table, and in the **Monitor > Alerts** view.

---

**Procedure**

1. From the **System Settings** menu in the top right corner, select **System Settings**.  
The **System Settings** window appears, showing approved and pending certificates.
2. Scroll to connections that are **Pending Approval**, and expand the rows.
3. For each one, scroll to the bottom of the information about the required certificate, and click **Confirm**.

## 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, Frontend internal views, and Property Sheet
- Log level

**Procedure**

1. From any location in the GUI, open the **System Settings** 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 52** User Preferences window

The screenshot shows the 'User Preferences' window with the following settings:



- General**: Refresh data every  seconds
- Login**: ☐ Clear host history
- Dashboard I/O Workload**: Average calculation will include the last  seconds; ☐ Show advanced dashboard
- System Clock**: ☐ Show system clock
- Property Sheet**: ☒ Show property sheet in advanced mode
- Frontend internals**: ☐ Show WOLs
- Backend internals**: ☒ Show internal backend views
- Support**: Log level: ; ☒ Pop up in case of uncaught exception

Buttons at the bottom: **Apply** and **Close**.

**Table 18** User Preferences

Item	Description
General: Refresh data every $n$ seconds	Controls the rate at which data displayed in the GUI is refreshed, in seconds (Default: 10 seconds)

**Table 18** User Preferences (continued)

Item	Description
	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.
Login: 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 <i>n</i> seconds	Controls the time period used when averages are computed and displayed by the GUI (default: 10 seconds)
Show advanced dashboard	<p>When selected, (default), includes more details on some tiles in the Dashboard view.</p> <p>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.</p> <p>The  symbol means that the number displayed is the average taken during the last <i>n</i> seconds. <i>n</i> can be configured in <b>Dashboard I/O Workload</b> in this window.</p> <p>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 <b>Dashboard I/O Workload</b> in this window.</p>
System Clock	Show system clock on the Dashboard
Show Property Sheet in advanced mode	<p>Displays additional details in Property Sheets:</p> <ul style="list-style-type: none"> <li>Capacity section—Snapshot Capacity Reserved</li> <li>Rebuild/Rebalance—Data Movement Jobs</li> <li>RAM Read Cache—Cache Evictions, Cache Entry, and Cache Skip tables</li> </ul> <p>These details are usually only relevant for advanced users and technical support purposes.</p>
Frontend Internals: Show VVols	Displays additional information for VVols
Backend Internals: Show internal Backend views	Displays additional options for Backend table views. These options are recommended only for advanced users and technical support purposes.
Support: 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.

**Table 18** User Preferences (continued)

Item	Description
	<div>Note</div> <div>Trace and Debug options may affect system performance, and are usually only recommended for technical support purposes.</div> <div>For more information about logs, see your system's Log Collection Technical Notes.</div>

# CHAPTER 12

## Using the VMware Plug-in

The following topics describe how to use the VMware plug-in (the “plug-in”) to view and provision ScaleIO components.

- [VMware Plug-in overview](#)..... 206
- [Configuring components](#)..... 207
- [Viewing components](#).....216

## 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:








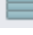
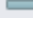

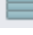




- Basic: The basic configurations are all performed the same way. The process is described just once.
- Advanced: Each advanced configuration setting has a unique dialog box, which is described in [“Configuring components—advanced”](#).

The following table lists the activities you can perform and categorizes each as basic or advanced:

**Table 19** 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 19** Plug-in activity matrix (continued)

Object	Perform this activity	Basic or advanced	Access from this screen
	Configure RAM Read Cache	Basic <sup>a</sup>	 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 <i>ScaleIO Deployment Guide</i> .	 EMC ScaleIO
	Upgrade SDC	Advanced	 SDCs
	Update SDC Parameters	Advanced. See the <i>ScaleIO Deployment Guide</i> .	 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 <sup>a</sup>	 Volumes
Fault Set	Create a Fault Set	Basic <sup>b</sup>	 Protection Domains
Device	Clear a device error  <b>Note</b> 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 19** 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 35. Failure to do so may prevent creation of volumes.

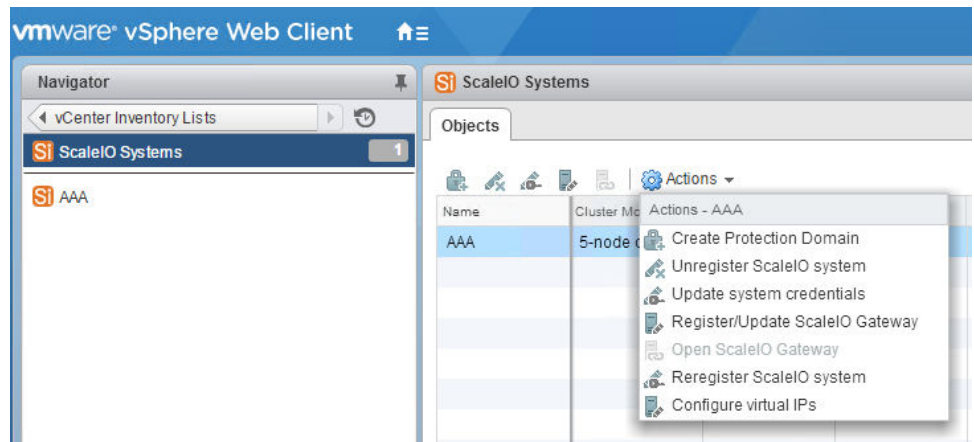
## Configuring components—basic

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 33.

---

## 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

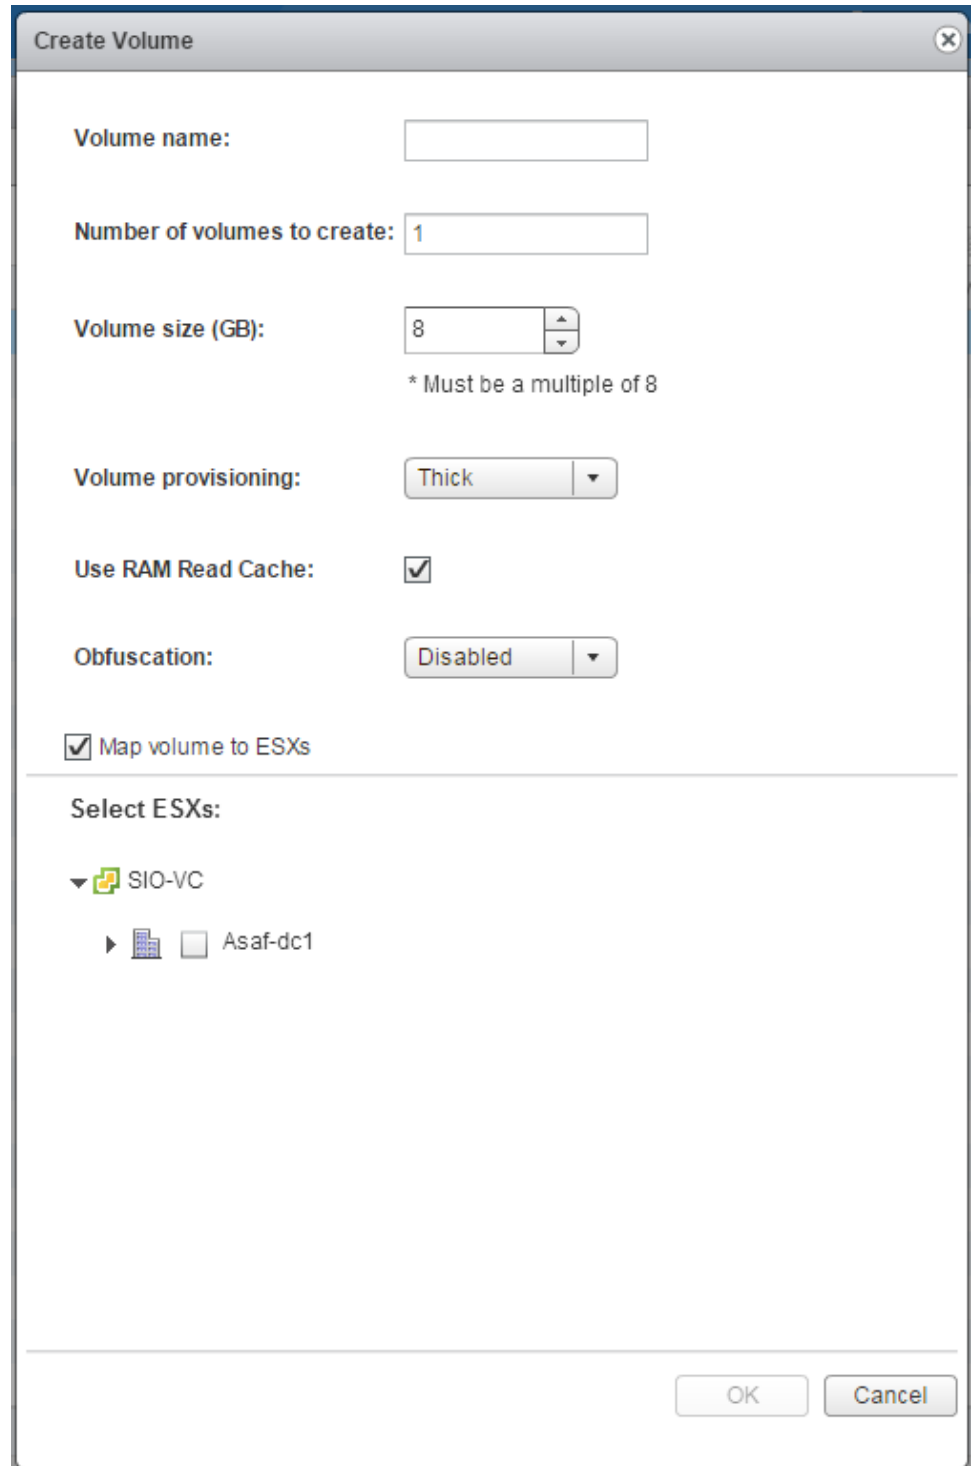
This section describes how to use the plug-in to create, map, and unmap volumes in the VMware environment. You can map volumes to SDCs in the same step, or you can map the volume after it has been created.

#### 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:



The image shows a 'Create Volume' dialog box with the following fields and options:

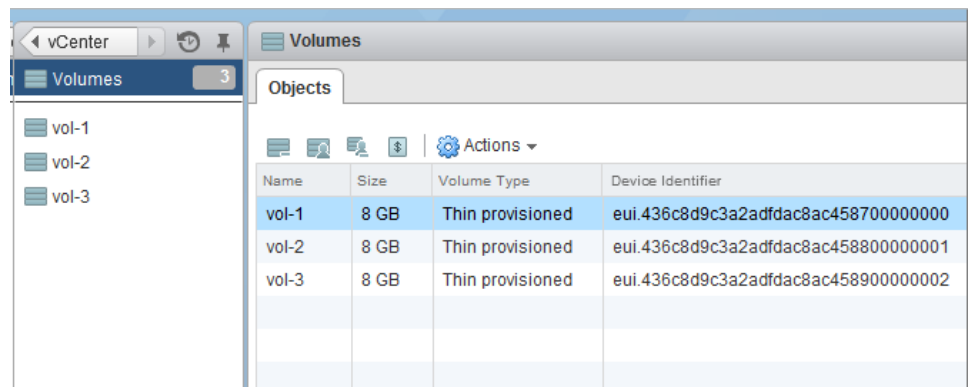
- Volume name:** A text input field.
- Number of volumes to create:** A text input field containing the value '1'.
- Volume size (GB):** A numeric input field with a spinner, showing '8'. Below it is a note: '\* Must be a multiple of 8'.
- Volume provisioning:** A dropdown menu set to 'Thick'.
- Use RAM Read Cache:** A checked checkbox.
- Obfuscation:** A dropdown menu set to 'Disabled'.
- ☒ **Map volume to ESXs**
- Select ESXs:** A section with a tree view showing 'SIO-VC' expanded, with a sub-entry 'Asaf-dc1' that has an unchecked checkbox next to it.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom right.

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:



Name	Size	Volume Type	Device Identifier
vol-1	8 GB	Thin provisioned	eui.436c8d9c3a2adfdac8ac458700000000
vol-2	8 GB	Thin provisioned	eui.436c8d9c3a2adfdac8ac458800000001
vol-3	8 GB	Thin provisioned	eui.436c8d9c3a2adfdac8ac4589000000002

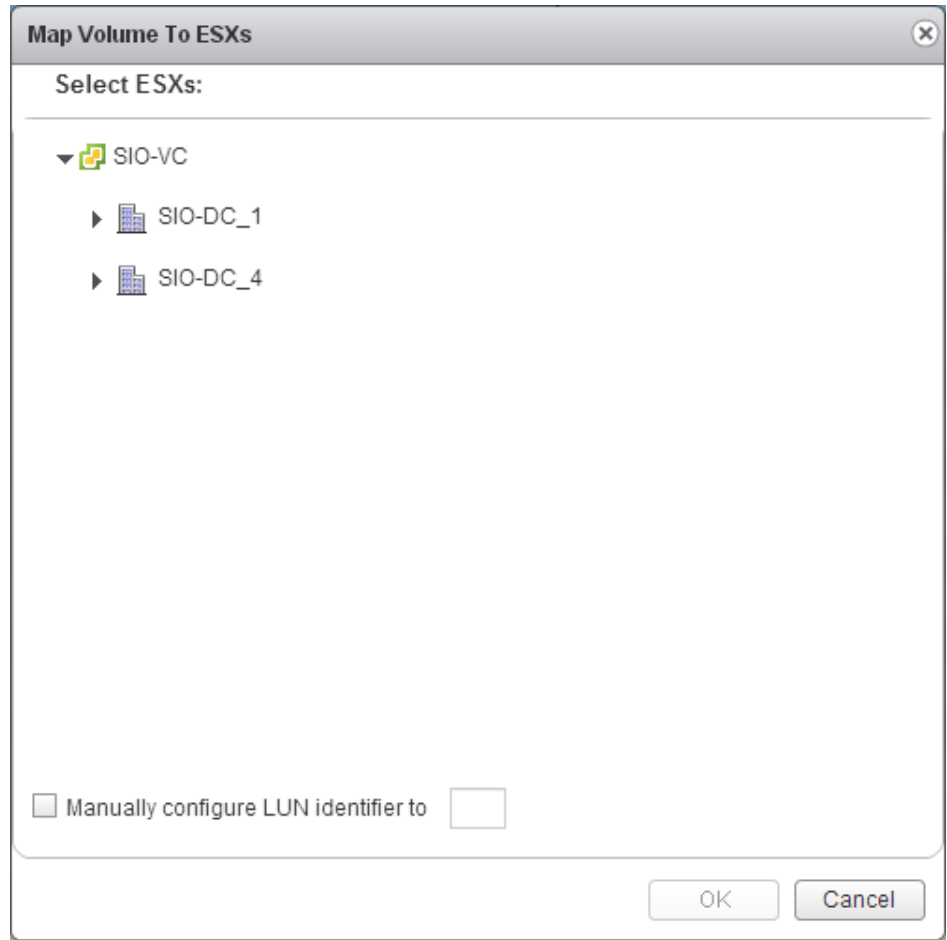
## Mapping volumes

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**.

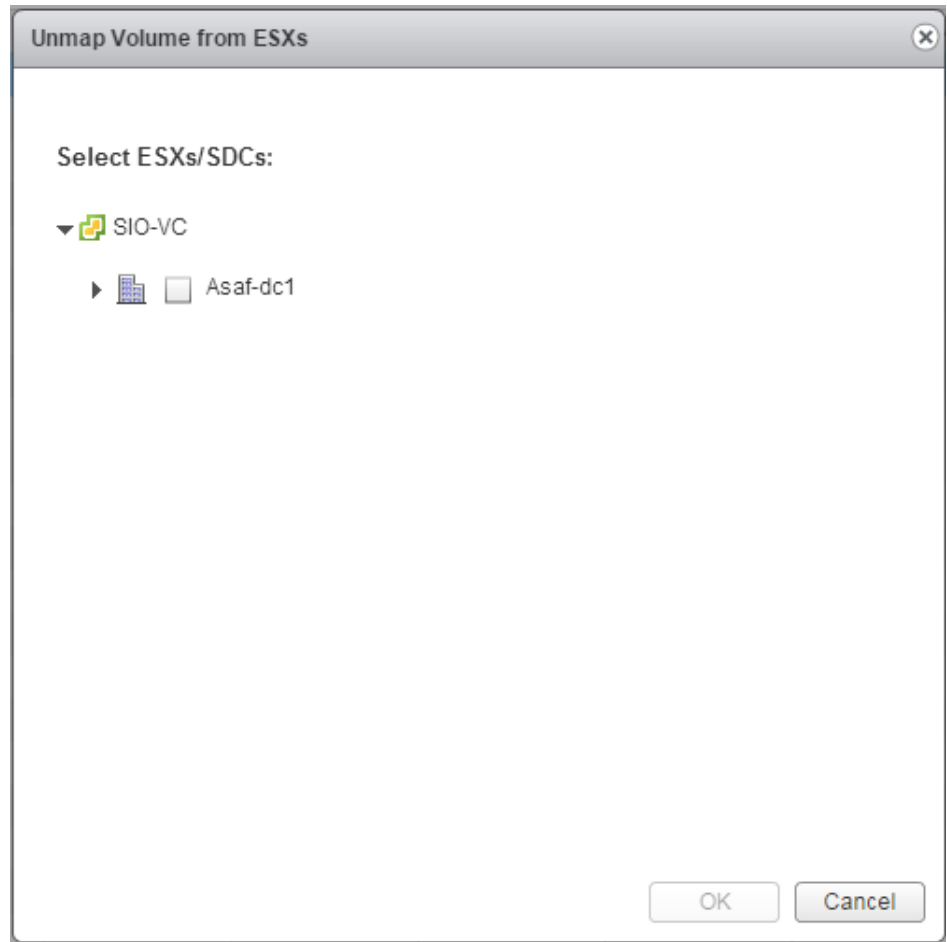
## Unmapping a volume

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.

---

### 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*.

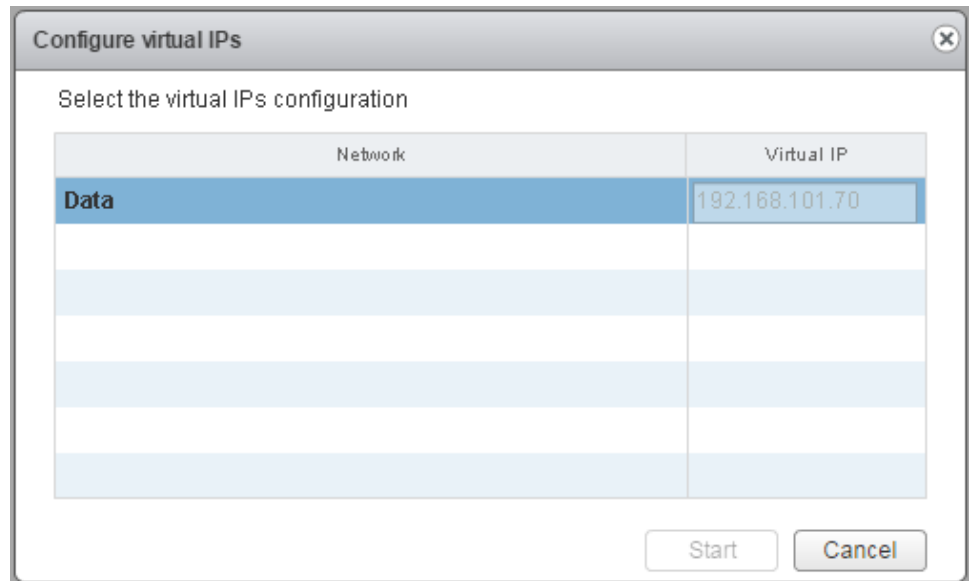
Note that there are differences in the way to perform this task, depending on the SDC version, as described in the guide.

## 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 53** 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 215.

## 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:

Name	Cluster Mode	Cluster state	Clients Number	Connected Client...	Servers Number	Devices Number	Vol
10.76.60.212	Cluster Mode	Normal	3	3	2	2	0

To drill down for more details, double-click the displayed details:



**ScaleIO Systems** | **sio9** | Actions

**Summary** | Related Objects

**ScaleIO Gateway**

- System name: 10.76.60.212
- Cluster mode: Cluster Mode
- Cluster state: Normal
- IP address: 10.76.60.212
- SVM Name: ScaleIO-10.103.111.8-GW
- ESX Name: 10.103.110.9

**Primary MDM**

- IP Addresses: 10.76.60.212
- Port: 9011
- SVM Name: ScaleIO-10.103.111.10
- ESX Name: 10.103.110.10

**Secondary MDM**

- IP Addresses: 10.76.60.211
- Port: 9011
- SVM Name: ScaleIO-10.103.111.9
- ESX Name: 10.103.110.9

**Tiebreaker MDM**

- IP Addresses: 10.76.60.213
- Port: 9011
- SVM Name: ScaleIO-10.103.111.11
- ESX Name: 10.103.110.11

**Management Network Configurations**

- IP Addresses: 10.103.111.9, 10.103.111.10
- Port: 6611

When you select **Storage Pools**, the following details appear:

Name	System Name	Total Capacity Li...	Devices	Volumes	RAM Read Cache	Write Handling M...	Zero Padding Policy
default	sio9	0.0 GB	0	0	Enabled	Cached	Disabled
sp	sio9	21.8 TB	24	1	Enabled	Cached	Disabled

When you drill-down on this screen, the following details appear:


**Storage Pools** | **default** | Actions

**Related Objects**

Protection Domains | Volumes | Devices

System Name	State	Total Capacity Li...	Storage Pools	Fault Sets	SDSs	Devices
pd	Active	407 GB	2	0	2	2

You can view the properties of all the ScaleIO components in the menu:



# PART 4

## Reference

The chapters in this part of the guide describe various topics related to advanced management.

Chapters include:

[Chapter 13, "Common Tasks"](#)

[Chapter 14, "System events"](#)

[Chapter 15, "ScaleIO on Xen"](#)

[Chapter 16, "Configuring ScaleIO in OpenStack Environments"](#)

[Chapter 17, "SNMP Trap Support"](#)

[Chapter 18, "ScaleIO SDC on AIX server"](#)



# CHAPTER 13

## Common Tasks

The following topics describe common tasks that are performed when working with ScaleIO.

• <a href="#">Install the ScaleIO GUI</a> .....	222
• <a href="#">Log in to the ScaleIO GUI</a> .....	222
• <a href="#">Connection and disconnection information</a> .....	223
• <a href="#">Add LIA to a system to enable automated upgrade</a> .....	223
• <a href="#">Associating ScaleIO volumes with physical disks</a> .....	224
• <a href="#">Port usage and changing default ports</a> .....	227
• <a href="#">Adding an external SDC to an existing ScaleIO system</a> .....	228
• <a href="#">Changing the LIA configuration file</a> .....	231
• <a href="#">Cleaning the ScaleIO VMware environment and performing a clean install</a> .....	231
• <a href="#">Configuring ScaleIO devices in Linux LVM</a> .....	233
• <a href="#">Configuring session timeout parameters</a> .....	234
• <a href="#">Fixing keytool errors</a> .....	234
• <a href="#">Installing Java on SUSE 12 servers</a> .....	235
• <a href="#">SVM manual memory allocation</a> .....	235
• <a href="#">Upgrading Java</a> .....	237
• <a href="#">Mounting ScaleIO</a> .....	237
• <a href="#">The ScaleIO Gateway web server isn't responding</a> .....	238
• <a href="#">Upgrading the Gateway when a custom certificate is used</a> .....	240
• <a href="#">Uploading a new OVA</a> .....	240
• <a href="#">Using the same data network for different NICs</a> .....	241
• <a href="#">What to do when the default self-signed certificate expires</a> .....	241
• <a href="#">Add another IP address subnet to an MDM cluster</a> .....	241
• <a href="#">Shutdown or restart a node gracefully</a> .....	243
• <a href="#">Deployment of ScaleIO using a non-root user</a> .....	246

## 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 either from the product ISO or the [EMC Support Site](#).

### Procedure

1. Install the GUI:

- Windows:

```
EMC-ScaleIO-gui-2.5-<build>.X.msi
```

- Linux:

```
rpm -i EMC-ScaleIO-gui-2.5-<build>.X.noarch.rpm
```

- Debian (run with administrator privileges):

```
sudo dpkg -i EMC-ScaleIO-gui-2.5-<build>.X.deb
```

### After you finish

To log in to the GUI, see "Log in to the ScaleIO GUI."

## 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 "Install 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:

- Linux: Run the script `/opt/emc/scaleio/gui/run.sh`.
- Windows: Click **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.

## Add LIA to a system to enable automated upgrade

Add the LIA, a component that is required to use the Installation Manager to upgrade ScaleIO physical server system components.

### Before you begin

To determine if the LIA is installed, run the following command on any server in the system:

```
rpm -qa | grep -i LIA
```

If LIA is not installed, you must install it before performing the upgrade.

Physical machine upgrade uses the Installation Manager (IM, part of the ScaleIO Gateway), together with the LIA of the new version, to orchestrate the upgrade.

### Procedure

1. Install the LIA component on every node, by running the following command:

```
TOKEN=<LIA_password> rpm -i <full rpm path to LIA file>
```

Example:

```
TOKEN=Scaleio123 rpm -i EMC-ScaleIO-lia-2.5-  
<build>.X.<flavor>.x86_64.rpm
```

The password must meet the following criteria:

- Between 6 and 31, ASCII-printable characters
- No blank spaces

- Include at least 3 of the following groups: [a-z], [A-Z], [0-9], special chars (!@#\$ ...)
2. Import the system installation ID into the LIA:
    - a. Create the following file:  
`/opt/emc/scaleio/lia/cfg/installation_id.txt`
    - b. Query the MDM for the installation ID by running the following command:
 

```
scli --query_all|grep "Installation ID"
```
    - c. Copy the installation ID into the new file.
    - d. Restart the LIA service by running the following command:
 

```
pkill lia
```
  3. Repeat the previous steps on every node in the system.

### Results

LIA is installed.

## 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 ad99eaab00000000 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 HarddiskX` 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\sdcb\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**

## Volume information - AIX

On AIX servers, associate the ScaleIO volume ID with the AIX physical device.

Retrieve the CuAt volume value:

#### Procedure

1. On the SDC host, run the following command to get the operating system volume information:

```
#odmget -q "name like scinid* and attribute=vol_id" CuAt
```

Output, similar to the following, is displayed:

```
CuAt:
name = "scinid0"
attribute = "vol_id"
value = "e120a92d00000000"
type = "R"
generic = "D"
rep = "s"
nls_index = 22
[root@cnode02 /]#odmget -q "name like scinid* and
attribute=vol_id" CuAt

CuAt:
name = "scinid2"
attribute = "vol_id"
value = "e120a92f00000002"
type = "R"
generic = "D"
rep = "s"
nls_index = 22

CuAt:
name = "scinid8"
attribute = "vol_id"
value = "e120a93500000008"
type = "R"
generic = "D"
rep = "s"
nls_index = 22
```

```

CuAt:
name = "scinid0"
attribute = "vol_id"
value = "e120a92d00000000"
type = "R"
generic = "D"
rep = "s"
nls_index = 22

```

You can get information for a single volume, by using this command:

```
#odmget -q "name=scinid0 and attribute=vol_id" CuAt
```

2. Match the value of the `value` field with the ScaleIO volume ID.

## 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 20** 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	<b>Note</b> 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: <ul style="list-style-type: none"> <li>Linux: Run <code>service scaleio-gateway restart</code></li> <li>Windows: Restart the EMC ScaleIO Gateway service</li> </ul>
Gateway-Installation	443 (or 8443,	REST over HTTPS	<gateway installation	ssl.port=443 (or 8443)	

**Table 20** Default ports (continued)

Port used by	Port #	Protocol	File to change	Field to modify (or to add, if it does not exist)	Notes
Manager/REST (secure)	together with 8080)		directory>/conf/catalina.properties		
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 operating system-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. On most servers (with the exception of hLinux), you can install the external SDC using the Installation Manager. For more information, see "Extending an existing ScaleIO system" in the *ScaleIO Deployment Guide*.

## Note

External SDC on RHEL 7.4 is supported on bare-metal servers only, not as guests on a hypervisor.

## Procedure

### 1. Install the SDC:

- RHEL/CentOS /Oracle Linux

```
MDM_IP=<LIST_VIP_MDM_IPS> rpm -i <SDC_PATH>.rpm
```

- 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.

## Install SDC on an AIX server and connect 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 an AIX server. The Installation Manager cannot be used.

### Procedure

1. Install the SDC:

```
MDM_IP=<LIST_VIP_MDM_IPS> rpm -i <SDC_PATH>.rpm
```

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. The SDC package is in a format similar to this: `EMC-ScaleIO-sdc-<version>-X.<build>.aix7.aix7.2.ppc.rpm`

### Results

The SDC is installed on the AIX 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
```

## 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.

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	<ul style="list-style-type: none"> <li>350 MB</li> </ul>		
MDM (Master/Slave)	<ul style="list-style-type: none"> <li><math>470 \text{ MB} + (500 \text{ KB} * 8 \text{ TB of volume capacity}) + (1.44 \text{ KB} * \text{number of volumes}) + (4 \text{ KB} * \text{number of SDS devices})</math></li> <li>Maximum supported volumes: 256 K</li> </ul>		
Tie Breaker MDM	<ul style="list-style-type: none"> <li>50 MB</li> </ul>		
SDS	<ul style="list-style-type: none"> <li><math>(\text{Base}) 536 \text{ MB} + (\text{RmCache Size}) * 1.15 + (\text{Storage capacity in TB}) * 53 \text{ MB}</math></li> <li>For SDS high performance profile, we add 195 MB.</li> </ul>		
SDC	<ul style="list-style-type: none"> <li><math>132 \text{ KB} + 23 \text{ MB} * (\text{number of MDMs}) + 25 \text{ KB} * (\text{number of SDSs}) + 1.5 \text{ KB} * (\text{number of volumes}) + 16 \text{ B} * (\text{number of volume blocks}) + 24 \text{ KB} * (8 \text{ TB of volume capacity})</math></li> <li>Volume blocks: 1 GB storage = 8 volume blocks</li> </ul>		
RFcache	<ul style="list-style-type: none"> <li><math>16 * (\text{cache\_size}/\text{page\_size})</math></li> <li>Commonly-used sizes:</li> </ul>		
	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

Component	Memory allocation rules		
	4 K	3.2 GB	6.4 GB

## Upgrading Java

Before changing the Java version of a node that is running the Gateway or AMS of ScaleIO v2.5 or later, you must prepare lockbox-related files.

The lockbox in ScaleIO v2.5 saves files in the Java folder of the ScaleIO Gateway and the AMS. These files need to be saved before any Java version update, then pasted back into the folder.

### Procedure

1. From the `jre\lib\ext` (or `jre/lib/ext` for Windows) Java folder, copy these files to a different folder:
  - `commons-lang3-3.6.jar`
  - `cryptoj-6.2.3.jar`
2. Update the Java version.
3. Paste these files back to the folder from where you copied them.

## 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 224.

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 <code>gatewayUser.properties</code> file
Windows, 64-bit	<code>C:\Program Files\EMC\ScaleIO\Gateway\webapps\ROOT\WEB-INF\classes\</code>
Linux	<code>/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes</code>

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-2.5-<build>.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
```

## Shutdown or restart a node gracefully

When performing tasks on a node that require it to be shutdown or restarted, do so gracefully.

Operating system upgrades and patches, as well as other maintenance activities, like part replacement, require shutting down or rebooting a node.

### Gracefully shut down or reboot a node

Prepare the server for a patching or maintenance operation (such as a part replacement) by entering the node into maintenance mode and shutting down/rebooting the node in a graceful fashion.

#### Before you begin

Ensure that you have admin rights for accessing the ScaleIO GUI. If necessary, the customer can give you the credentials.

#### Procedure

1. When shutting down/rebooting a node that is a Master MDM, it is recommended that you manually switch MDM ownership to a different node:
  - a. From the ScaleIO CLI (SCLI), run:

```
scli --query_cluster
```

---

#### Note

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

- ESXi (SVM) — `scli`
  - Linux — `scli`
  - Windows — `C:\Program Files\emc\scaleio\MDM\bin`
- 

- b. If the node's IP addresses are included in the `--query_cluster` output, the faulty node has a role of either MDM or TieBreaker (TB), in addition to its SDS role.

If the node's IP address is located in the Master MDM role, a switch-over action is required.

## c. Switch MDM ownership to a different node:

```
scli -switch_mdm_ownership (-new_master_mdm_id <ID> | --
new_master_mdm_ip <IP> | --new_master_mdm_name <NAME>)
```

The node remains in the cluster. The cluster will be in degraded mode after it is powered off, until the faulty component or patch operation in the node is fixed and the node is powered back on.

## d. Verify that the cluster status shows that the node is not the Master MDM anymore:

```
scli --query_cluster
```

Output similar to the following should appear, with the relevant node configuration and IP addresses for your deployment:

```
Cluster:
  Mode: 5_node, State: Normal, Active: 5/5, Replicas: 3/3
  Virtual IP Addresses: 9.20.10.100, 9.20.110.100
Master MDM:
  ID: 0x775afb2a65ef1f02
  IP Addresses: 9.20.10.104, 9.20.110.104, Management
IP Addresses: 10.136.215.239, Port: 9011, Virtual IP
interfaces: sio_d_1, sio_d_2
  Version: 2.0.13000
Slave MDMs:
  ID: 0x5b2e9f273b7af9b0
  IP Addresses: 9.20.10.105, 9.20.110.105, Management
IP Addresses: 10.136.215.223, Port: 9011, Virtual IP
interfaces: sio_d_1, sio_d_2
  Status: Normal, Version: 2.0.13000
  ID: 0x5828f65b15e778f1
  IP Addresses: 9.20.10.102, 9.20.110.102, Management
IP Addresses: 10.136.215.232, Port: 9011, Virtual IP
interfaces: sio_d_1, sio_d_2
  Status: Normal, Version: 2.0.13000
Tie-Breakers:
  ID: 0x6618e0b804644ca4
  IP Addresses: 9.20.10.101, 9.20.110.101, Port: 9011
  Status: Normal, Version: 2.0.13000
  ID: 0x12534ccb3d28fee3
  IP Addresses: 9.20.10.103, 9.20.110.103, Port: 9011
  Status: Normal, Version: 2.0.13000
```

In the example output, the Master MDM IP addresses are:

```
IP Addresses: 9.20.10.104, 9.20.110.104, Management IP
Addresses: 10.136.215.239
```

The Slave IP addresses are:

```
IP Addresses: 9.20.10.105, 9.20.110.105, Management IP
Addresses: 10.136.215.223
IP Addresses: 9.20.10.102, 9.20.110.102, Management IP
Addresses: 10.136.215.232
```

## 2. Move all applications to a different node:

- On an ESXi node that is not a cluster member, and that is not configured for HA and DRS, migrate the VMs to another ESXi.

- On a Linux or a Windows node, migrate the applications (or the VMs, if the node is running a hypervisor).

---

#### Note

In non-hypervisor environments, ask the customer for assistance in moving applications from the node.

---

3. Log in to the ScaleIO GUI as an admin user.
4. In ScaleIO **Backend** view, select **By SDSs** table view.
5. Right-click the SDS node you are shutting down/rebooting, and select **Enter Maintenance Mode**.
6. In the **Enter maintenance mode** window, wait for rebalance operations to finish, ensure that there are no errors, and then click **OK**.
7. When the operation finishes successfully, click **Close**.  
The node's IP address appears with a wrench next to it.
8. On an ESXi node:
  - a. Log in to the vCenter via the vSphere Web Client, and locate the relevant ESXi IP address.
  - b. Select the SVM, and from the **Actions** > **Power** menu, select **Shut Down Guest OS**.
  - c. When the SVM is off, right-click the ESXi node and select **Enter Maintenance Mode**.
9. If you are applying a patch:
  - a. Run the patch.
  - b. Reboot the node, if necessary.
10. For part replacement or to shut down/reboot a node:
  - a. Obtain customer permission to shut down the node.
11. Gracefully shut down/reboot the node using the relevant API for the operating system.

---

#### Note

On a Linux or Windows node, no checks are required for a graceful shutdown after entering the SDS into maintenance mode.

---

## Return the node to operation

To return the node to operation, perform the following steps:

### Procedure

1. Power on the node, or if rebooting, wait for the node to start booting.  
The OS will boot up for Windows and Linux operating systems. For Windows and Linux nodes, all ScaleIO processes will start up automatically.
2. For an ESXi node, perform the following:
  - a. From the vSphere Web Client, ensure that the node is displayed as on and connected in both **Hosts** and **Clusters** view.

- b. Right-click the node and select **Exit Maintenance Mode**.
    - c. Expand the server and select the ScaleIO VM. If the SVM does not power on automatically, power it on manually.
  3. After the node is up, perform the following checks in the ScaleIO GUI:
    - a. In the **Monitor > Alerts** view, make sure that no SDS disconnect message appears.
    - b. If the node was an MDM cluster member, in the Dashboard **Management** tile, verify that the cluster is no longer degraded.
    - c. In the **Frontend** tab > **SDCs** view, check the SDC to which the node IP is assigned, and make sure that it is connected.
  4. In the ScaleIO GUI **Backend** view, in **By SDSs** table view, right-click the SDS and select **Exit Maintenance Mode**.
  5. In the **Action** window, click **OK**.
  6. Wait for the rebalance operations to finish.
- The node is now operational and application I/O can be started on the node. For ESXi nodes, you can migrate VMs to the node.

## Deployment of ScaleIO using a non-root user

ScaleIO can be deployed or extended in Linux environments using a non-root sudo user in non-interactive mode.

Sudo is a program that allows a user to run or install a program as the root user. A sudo user can be created to deploy ScaleIO.

In order to successfully deploy or extend ScaleIO with a non-root user, the non-root user must meet the following conditions:

- The username included in the CSV file must already exist.
- The non-root user must be a sudo user.
- The non-root user must be in non-interactive mode.
- The requirement for TTY must be disabled.

In the CSV file used for deployment, you must indicate that you are intending to use a sudo non-root username by appending the string "(sudo)" to the user name in the Username field. For example, if you are using a non-root user with the username "non\_root", enter the string "non\_root(sudo)" in the username field of the CSV file.

## Configure a non-root non-interactive sudo user

In Linux, you can deploy or extend ScaleIO with a non-root user. You must configure a non-root sudo user in non-interactive mode.

### Before you begin

The following procedure details one method for configuring a non-root non-interactive sudo user. Perform the commands from the operating system console of where you want the gateway to deploy the ScaleIO system.

### Procedure

1. Create a user group named "admin".

```
groupadd admin
```

2. Create a user named "non\_root" and add it to the admin group.

```
useradd -G admin non_root
```

3. Change the password of the non\_root user.

```
passwd non_root
```

When prompted, enter the new password and then confirm it by entering it again.

4. Open the sudoers /etc/sudoers file for editing.

```
visudo
```

5. Search the sudoers file for "## Same thing without a password".

```
:s/## Same thing without a password
```

6. In the line below the search result, add the text "%admin ALL=(ALL) NOPASSWD: ALL" and then exit the vi editor.

Type the following command to exit: :q

7. Create a hidden directory in the non\_root user's home directory to store the SSH configuration.

```
mkdir /home/non_root/.ssh
```

8. Copy the SSH configuration from the root user to the non\_root user's directory.

```
cp -rf /root/.ssh/* /home/non_root/.ssh/
```

9. Open the sudoers /etc/sudoers file for editing.

```
visudo
```

10. Search the sudoers file for "Defaults requiretty" and replace it with "Defaults ! requiretty".





# CHAPTER 14

## System events

The following topics describe ScaleIO system events and alerts.

- [System events overview](#) ..... 250
- [Event format](#).....250
- [Viewing events](#)..... 251
- [Event list](#).....254

## 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](#) on page 96.

### Description/Notes

Displays events, which can be filtered by optional switches.

### Parameters

Parameter	Description
Options:	
<code>--min_severity &lt;SEVERITY&gt;</code>	Displays events with at least the specified minimum severity
<code>--severity &lt;SEVERITY&gt;</code>	Displays events with the specified severity
<code>--from_id &lt;ID&gt;</code>	Displays all events starting from the given ID
<code>--to_id &lt;ID&gt;</code>	Displays all events ending at the given ID
<code>--from_date &lt;ID&gt;</code>	Displays all events starting from the given date
<code>--to_date &lt;ID&gt;</code>	Displays all events ending at the given date
<code>--grep &lt;TEXT&gt;</code>	Displays events containing the specified text
<code>--full</code>	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 15

## ScaleIO on Xen

The following topics describe using ScaleIO on Xen.

- [Overview of ScaleIO on Xen](#).....268
- [Adding a volume in XenServer environment](#)..... 268
- [Removing a ScaleIO volume from Xen](#).....269
- [Modifying the size of a ScaleIO volume](#).....270
- [Xen v6.5 High Availability](#).....271

## 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=lvnobra
```

---

**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 PBD 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 16

## Configuring ScaleIO in OpenStack Environments

The following topics contain information about ScaleIO provisioning in an OpenStack cloud operating system environment.

- [Overview](#) ..... 274

## 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 ScaleIO.

For more information, refer to OpenStack documentation (<https://docs.openstack.org>); especially the *Cinder Configuration Guide*.

# CHAPTER 17

## SNMP Trap Support

The following topics describe ScaleIO support for SNMP.

- [General](#).....276
- [Supported alerts and event numbering conventions](#).....276
- [Configure SNMP properties after deployment](#).....303
- [ScaleIO.mib file](#).....305

## 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 303.

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 21** 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	SIO01.02.000001	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	SIO01.02.000002	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	SIO01.02.000003	2 (Warning)	Purchase a license and install it.
Oscillating failures reported	OBJECT_HAS_OSCILLATING_FAILURES	System.Oscillating_Failures.Object	SIO01.03.000001	2 (Warning)	Check oscillating failures of the component and take action

**Table 21** 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
		_has_oscillating_failures			accordingly. If the oscillating failure does not indicate a problem, change the settings of the oscillating failure window to suppress this alert
There are oscillating network failures	OBJECT_HAS_OSCILLATING_NETWORK_FAILURES	System.Oscillating_Failures.OBJECT_HAS_OSCILLATING_NETWORK_FAILURES	SIO01.03.000002	2 (Warning)	Check the oscillating failure report, that can be accessed from one of the management interfaces. Check whether there is a 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	SIO01.04.000001	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	SIO01.04.000002	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	SIO01.04.000004	5 (Critical)	Configure MDM credentials on the ScaleIO Gateway using the SioGWTool
System upgrade is in progress	UPGRADE_IN_PROGRESS	System.Upgrade.	SIO01.05.000001	3 (Error)	Monitor the upgrade process,

**Table 21** 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
		UPGRADE_IN_PROGRESS			and check that it is completed successfully.
ScaleIO Gateway version is too old	GW_TOO_OLD	System.Upgrade.GW_TOO_OLD	SIO01.05.000002	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	SIO02.01.000001	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	SIO02.01.000003	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	SIO02.01.000004	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.
Backward rebuild cannot proceed	BKWD_REBUILD_STUCK	MDM.MDM_Cluster.BKWD_REBUILD_STUCK	SIO02.01.000005	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 21** 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
Rebalance cannot proceed	REBALANCE_STUCK	MDM.MDM_Cluster.REBALANCE_STUCK	SIO02.01.000006	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	SIO02.01.000007	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.
Cannot connect to the MDM cluster, but the cluster itself is operational	MDM_CONNECTION_LOST	MDM.MDM_Cluster.MDM_CONNECTION_LOST	SIO02.01.000008		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	SIO02.01.000009	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.



**Table 21** 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
Inactive Protection Domain	PD_INACTIVE	MDM.Protection_Domain.Protection_Inactive	SIO02.02.000001	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.
Storage Pool has failed capacity	STORAGE_POOL_HAS_FAILED_CAPACITY	MDM.Storage_Pool.Storage_Pool_has_Failed_Capacity	SIO02.04.000001	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	SIO02.04.000002	3 (Error)	For the given Storage Pool, for some blocks, one of the two copies is inaccessible. 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	SIO02.04.000003	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

**Table 21** 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
					possible, or add physical storage.
Capacity utilization above high threshold	CAPACITY_UTILIZATION_ABOVE_HIGH_THRESHOLD	MDM.Storage_Pool.Capacity_Utilization_Above_High_Threshold	SIO02.04.00 00004	2 (Warning) 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.
Failure recovery capacity is below the threshold	FAILURE_RECOVERY_CAPACITY_BELOW_THRESHOLD	MDM.Storage_Pool.Failure_Recovery_Capacity_Below_Threshold	SIO02.04.00 00005	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_Than_Largest_Fault_Unit	SIO02.04.00 00008	2 (Warning)	Increase the "spare percentage" configured in the system for the Storage Pool, so that the capacity 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	SIO02.04.00 00009	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.

**Table 21** 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 does not meet the minimum requirement of 3 fault units	NOT_ENOUGH_FAULT_UNITS_IN_SP	MDM.Storage_Pool.Not_Enough_Fault_Units	SIO02.04.000010	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 > Connection.	UNTRUSTED_CERTIFICATE	MDM.CERTIFICATE.UNTRUSTED_CERTIFICATE	SIO02.05.000001	3 (Error)	Pending for approval certificates can be viewed and approved via System Settings > Connection.
Master MDM Certificate is about to expire	CERTIFICATE_ABOUT_TO_EXPIRE	MDM.CERTIFICATE.CERTIFICATE_ABOUT_TO_EXPIRE	SIO02.05.000002	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	SIO02.05.000003	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	SIO02.05.000004	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_SELF_SIGNED_CERTIFICATE_NOT_TRUSTED	SIO02.05.000005	5 (Critical)	Check the certificate, and trust it if you see fit to do so

**Table 21** 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
MDM does not support secure connections	MDM_SECURE_CONNECTION_NOT_SUPPORTED	MDM.CERTIFICATE.MDM_SECURE_CONNECTION_NOT_SUPPORTED	SIO02.05.000006	5 (Critical)	Check MDM cluster nodes
The validity period of the certificate presented by the Master MDM starts in the future	MDM_CERTIFICATE_NOT_YET_VALID	MDM.CERTIFICATE.MDM_CERTIFICATE_NOT_YET_VALID	SIO02.05.000007	5 (Critical)	The time and date on the computer where the certificate was created is 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	SIO02.05.000008	5 (Critical)	Trust the CA certificate if you see fit
SDS is disconnected	SDS_DISCONNECTED	SDS.SDS.SDS_Disconnected	SIO03.01.000001	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	SIO03.01.000002	3 (Error) 2 (Warning) according to disconnect count and hard-coded	The SDS connection is fluctuating due to an unstable network connection. Check

**Table 21** 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 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	SIO03.01.000003	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 configured RAM Read-Cache memory to match the allocation conditions.

**Table 21** 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
DRL mode: Hardened	DRL_MODE_NON_VOLATILE	SDS.SDS.DRL_MODE_NON_VOLATILE	SIO03.01.000004	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	SIO03.01.000005	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	SIO03.01.000006	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	SIO03.01.000007	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	SIO03.01.000008	2 (Warning)	There is not enough RAM available on the server for Read Flash Cache optimal operation. Increase the amount of available RAM.

**Table 21** 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 driver path is invalid	RFCACHE_INVALID_DRIVER_PATH	SDS.SDS.RFCACHE_INVALID_DRIVER_PATH	SIO03.01.000009	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 not consistent	RFCACHE_INCONSISTENT_SOURCE_CONFIGURATION	SDS.SDS.RFCACHE_INCONSISTENT_SOURCE_CONFIGURATION	SIO03.01.000010	2 (Warning)	Check RFcache state of all disks in the pool and adjust them so that all disks have the same caching state.
RFcache source configuration is not consistent	RFCACHE_INCONSISTENT_CACHE_CONFIGURATION	SDS.SDS.RFCACHE_INCONSISTENT_CACHE_CONFIGURATION	SIO03.01.000011	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	SIO03.01.000012	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_MISMATCH	SDS.SDS.RFCACHE_API_ERROR_MISMATCH	SIO03.01.000013	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.

**Table 21** 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
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 21** 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	SIO03.02.000005	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	SIO03.02.000006	3 (Error)	Check the device, and if necessary, replace it
Device malfunction has been detected.	DEVICE_ERROR_NOTICE	SDS.Device.Device_ErrorNotice	SIO03.02.000007	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	SIO03.02.000008	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	SIO03.02.000009	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	SIO03.02.000011	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	SIO03.02.000013	3 (Error)	Consider replacing the disk.
The SDC is either down or unreachable over the network	SDC_DISCONNECTED	SDC.SDC.SDC_DISCONNECTED	SIO04.01.000001	3 (Error)	Verify that the SDC service is up and running and that the network is

**Table 21** 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
					properly configured and connected.
No more SDCs can be defined on this system; the maximum has been reached	SDC_MAX_COUNT	SDC.SDC_Max_Count	SIO04.01.00 00002	3 (Error)	The maximum number of SDCs in the system has been reached (1024^2).
SSD lifespan is over endurance threshold	PHYSICAL_DRIVE_ENDURANCE_USED_ABOVE_THRESHOLD	Storage_Controller.Physical_Drive.Endurance_Used_Above_Threshold	SIO06.01.00 00004	5 (Critical)	Replace the SSD device.
The Physical Disk temperature has exceeded the configured threshold	PHYSICAL_DRIVE_INVALID_TEMPERATURE	Storage_Controller.Physical_Drive.INVALID_TEMPERATURE	SIO06.01.00 00005	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled. Ensure that the internal chassis fans are working, and that there is adequate airflow.
The Physical Drive's cache is not disabled	PHYSICAL_DRIVE_INVALID_PD_CACHE_POLICY	Storage_Controller.Physical_Drive.INVALID_PD_CACHE_POLICY	SIO06.01.00 00006	3 (Error)	Disable the Physical Drive physical cache. Note: This may cause DI issues in power cycle scenarios.
Device was requested to use as DAS Cache but it is not	DEVICE_SHOULD_USE_AS_DAS_CACHE_BUT_IT_IS_NOT	Storage_Controller.Physical_Drive.Device_Should_Use_As_Das_Cache_But_It_Is_Not	SIO06.01.00 00007	3 (Error)	
The Logical Disk read policy is not	LOGICAL_DISK_INVALID_READ_AHEAD_POLICY	Storage_Controller	SIO06.02.00 00001	3 (Error)	Access the storage controller and

**Table 21** 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
set to "Read-Ahead"		r.Logical_Disk.INVALID_READ_AHEAD_POLICY			change the Logical Disk's read policy to "Read Ahead".
The Logical Disk IO policy is not Configured properly (Write-Back)	LOGICAL_DISK_INVALID_WRITE_BACK_POLICY	Storage_Controller.Logical_Disk.INVALID_WRITE_BACK_POLICY	SIO06.02.00 00002	3 (Error)	Check the Raid controller battery status, a fault in the battery will impact the Write-Back functionality.
Logical Disk access mode is not set to "Read-Write"	LOGICAL_DISK_INVALID_ACCESS_MODE	Storage_Controller.Logical_Disk.INVALID_ACCESS_MODE	SIO06.02.00 00003	5 (Critical)	Access the storage controller and change the logical disk's access mode to "Read-Write".
The Logical Disk RAID level is not set to RAID0	LOGICAL_DISK_INVALID_RAID_LEVEL	Storage_Controller.Logical_Disk.INVALID_RAID_LEVEL	SIO06.02.00 00004	3 (Error)	The RAID type is set in the storage controller is incorrect. Access the storage controller, verify that the Logical Disk does not contain any data that is not backed-up, destroy the Logical Disk, and re-create the logical disk as RAID0.
Logical Disk caching policy is not set to "DirectIO"	LOGICAL_DISK_INVALID_CACHE_POLICY	Storage_Controller.Logical_Disk.INVALID_CACHE_POLICY	SIO06.02.00 00005	3 (Error)	Access the storage controller, and change the RAID0 caching policy for the Logical Disk to "DirectIO".

**Table 21** 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
Logical disk is no longer being cached	LOGICAL_DISK_NO_LONGER_CACHED	Storage_Controller.Logical_Disk.LOGICAL_DISK_NO_LONGER_CACHED	SIO06.02.000006	3 (Error)	Disable the Physical Drive physical cache. Note: This may cause DI issues in power cycle scenarios.
The state of the backup battery in the storage controller is not optimal	BATTERY_INVALID_STATE	Storage_Controller.Battery.Invalid_State	SIO06.03.000001	5 (Critical) 3 (Error)	The backup battery is not fully charged, but it will recharge itself while the storage controller is powered on.
The backup battery in the storage controller needs to be replaced	BATTERY_REPLACEMENT_REQUIRED	Storage_Controller.Battery.REPLACEMENT_REQUIRED	SIO06.03.000002	5 (Critical)	The backup battery in the storage controller may be nearing the end of its working life. Replace the battery.
The storage controller battery temperature is not within the configured threshold	BATTERY_INVALID_TEMPERATURE	Storage_Controller.Battery.INVALID_TEMPERATURE	SIO06.03.000003	3 (Error)	Check the temperature in the server and at the data center. Check if a fan alert is raised in the node. Check the battery and the RAID controller and replace faulty items.
There is no backup battery installed in the storage controller	BATTERY_NOT_PRESENT	Storage_Controller.Battery.NOT_PRESENT	SIO06.03.000004	5 (Critical)	Install a backup battery in the storage controller.
The storage controller contains an incompatible battery pack	BATTERY_INVALID_PACK_ENERGY	Storage_Controller.Battery.Invalid_Pack_Energy	SIO06.03.000005	3 (Error)	Replace the storage controller battery with one that is compatible with your controller's model.

**Table 21** 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
RAID controller's battery has invalid voltage state	BATTERY_INVALID_VOLTAGE	Storage_Controller.Battery.Invalid_Voltage	SIO06.03.00 00006	3 (Error)	Check the RAID controller's battery. It may need to be replaced.
Hypervisor boot drive state is not OK	BOOT_DRIVE_INVALID_STATE	Node.Boot_Drive.Invalid_State	SIO07.01.00 00001	3 (Error)	The drive on which the ESXi hypervisor is installed has reported an error. Replace the drive and reinstall the ESXi on the new drive.
The storage controller is not operating optimally	STORAGE_CONTROLLER_INVALID_STATE	Node.Storage_Controller.Invalid_State	SIO07.02.00 00001	5 (Critical)	The storage controller may be faulty, and should be replaced.
The storage controller temperature has exceeded the configured threshold	STORAGE_CONTROLLER_INVALID_TEMPERATURE	Node.Storage_Controller.Invalid_Temperature	SIO07.02.00 00002	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
CacheCade license is not installed	STORAGE_CONTROLLER_CACHECADE_NOT_LICENSED	Node.Storage_Controller.Cachecade_Not_Licensed	SIO07.02.00 00003	3 (Error)	Install a CacheCade license on the storage controller.
Some of the disk slots in the storage controller are empty	STORAGE_CONTROLLER_NOT_ALL_SLOTS_FULL	Node.Storage_Controller.Not_All_Slots_Full	SIO07.02.00 00004	3 (Error)	Add more disks if needed. If a disk was removed during a FRU (Field Replacement Unit) operation, insert the new disk into the chassis.
Unable to query RAID controller due to insufficient permissions	NODE_FAILED_TO_QUERY_RAID_CONTROLLER_INVALID_PERMISSIONS	Node.Storage_Controller.NODE_FAILED_TO_QUERY	SIO07.02.00 00005	5 (Critical)	Verify that the user name and password of the ScaleIO VM match the user name and password configured in the

**Table 21** 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
		Y_RAID_CONTROLLER_INVALID_PERMISSIONS			AMS (System Settings menu > System Settings > Security option). Resolve the differences and try again.
The physical CPU socket is not enabled	SOCKET_DISABLED	Node.Cpu_Socket.Disabled	SIO07.03.00 00001	3 (Error)	Enable the CPU socket in the server BIOS.
Not all CPU cores are enabled	SOCKET_NOT_ALL_CORES_ENABLED	Node.Cpu_Socket.NOT_ALL_CORES_ENABLED	SIO07.03.00 00002	3 (Error)	Change the physical CPU core configuration in the server BIOS.
CPU is not operating at full speed	SOCKET_SPEED_IS_NOT_MAX_SPEED	Node.Cpu_Socket.SPEED_IS_NOT_MAX_SPEED	SIO07.03.00 00003	3 (Error)	Set the CPU clock speed in the server's BIOS to its maximum speed.
CPU temperature is not within threshold	CPU_SOCKET_INVALID_TEMPERATURE	Node.Cpu_Socket.INVALID_TEMPERATURE	SIO07.03.00 00004	5 (Critical) 3 (Error)	The server's CPU temperature has exceeded the configured threshold. Make sure that the server is properly cooled and that the CPU and internal chassis fans are active.
The CPU Voltage Regulator's temperature is higher than the configured threshold	CPU_SOCKET_INVALID_VR_TEMPERATURE	Node.Cpu_Socket.INVALID_VR_TEMPERATURE	SIO07.03.00 00005	5 (Critical) 3 (Error)	Make sure that the server is properly cooled and that the CPU and internal chassis fans are active.
The CPU Voltage Regulator's voltage is not within the configured threshold	CPU_SOCKET_INVALID_VR_VOLTAGE	Node.Cpu_Socket.INVALID_VR_VOLTAGE	SIO07.03.00 00006	5 (Critical) 3 (Error)	1. Verify that the power supply is functioning correctly. 2. Try to replace a port in

**Table 21** 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 Power Distribution Unit, or supply an external power source to check it. 3. Replace the power cable. 4. Replace the Power Supply Unit module.
CPU has invalid voltage state state	CPU_SOCKET_INVALID_VOLTAGE	Node.Cpu.Socket.Invalid_Voltage	SIO07.03.000007	5 (Critical)	Check the chassis power supply
RAM temperature is not within threshold	RAM_INVALID_TEMPERATURE	Node.Ram.Invalid_Temperature	SIO07.04.000001	5 (Critical) 3 (Error)	The temperature of one or more Server RAM modules exceeds the configured threshold. Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
RAM voltage regulator's temperature is not within threshold	RAM_INVALID_VR_TEMPERATURE	Node.Ram.INVALID_VR_TEMPERATURE	SIO07.04.000002	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional.
RAM voltage regulator's voltage is not within threshold	RAM_INVALID_VR_VOLTAGE	Node.Ram.INVALID_VR_VOLTAGE	SIO07.04.000003	5 (Critical) 3 (Error)	Ensure that the server and its environment are properly cooled and that all the chassis fans are functional. If the system still issues an alert, the DIMM may be faulty, and may have to be replaced.
The node temperature has	NODE_INVALID_TEMPERATURE	Node.Node.INVALID	SIO07.05.000001	5 (Critical) 3 (Error)	Ensure that the server and its

**Table 21** 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
exceeded the configured threshold		ID_TEMPERATURE			environment are properly cooled and that all the chassis fans are functional.
The voltage values on the node are not within the configured threshold	NODE_INVALID_VOLTAGE	Node.Node.INVALID_VOLTAGE	SIO07.05.000002	5 (Critical) 3 (Error)	1. Verify that the power supply is functioning correctly, and then try to replace a port in the Power Distribution Unit or check with external power. 2. Replace the power cable. 3. Replace the Power Supply Module.
Unable to connect to monitoring agent (BMC)	NODE_FAILED_TO_CONNECT_TO_BMC	Node.Node.FAILED_TO_CONNECT_TO_BMC	SIO07.05.000004	5 (Critical)	The BMC IP address or User credentials do not respond to AMS queries. Verify network connectivity from the AMS to the BMC IP address, and check that the BMC admin user and password have not been tampered with. Contact customer support for assistance if the problem persists.
SDC is not installed on this node	NODE_WITH_NO_SDC	Node.Node.NODE_WITH_NO_SDC	SIO07.05.000005	3 (Error)	Consider installing an SDC on this node, so that it can use ScaleIO volumes
Unable to connect to monitoring agent (SVM)	NODE_FAILED_TO_CONNECT_TO_SVM	Node.Node.FAILED_TO_CONNECT_TO_SVM	SIO07.05.000006	5 (Critical)	The SVM IP address or User credentials do not respond to AMS queries. Check the SVM state in vCenter and try to



**Table 21** 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
					open the SSH or the console to it. Then check the message log for errors.
Unable to connect to monitoring agent (ESX)	NODE_FAILED_TO_CONNECT_TO_ESX	Node.Node.FAILED_TO_CONNECT_TO_ESX	SIO07.05.000007	5 (Critical)	The ESX IP address or User credentials do not respond to AMS queries. Check the ESX state in vCenter and try to open the SSH or the console (via BMC http) to it. Then check the messages log for errors.
Unable to connect to monitoring agent (host)	NODE_FAILED_TO_CONNECT_TO_HOST	Node.Node.FAILED_TO_CONNECT_TO_HOST	SIO07.05.000007	5 (Critical)	The ESX IP address or User credentials do not respond to AMS queries. Check the ESX state in vCenter and try to open the SSH or the console (via BMC http) to it. Then check the messages log for errors.
Unable to connect to monitoring agent (vCenter)	NODE_FAILED_TO_CONNECT_TO_VCENTER	Node.Node.FAILED_TO_CONNECT_TO_VCENTER	SIO07.05.000008	5 (Critical)	Ensure that the vCenter configuration is correct, and ensure that the node's Mgmt port is routable to the vCenter IP address
Node serial number has changed	NODE_SERIAL_NUMBER_CHANGE	Node.Node.SERIAL_NUMBER_CHANGE	SIO07.05.000009	3 (Error)	The S/N of the Motherboard does not match the IP address of the ESX. Either the SATADOM was moved to a new server or the

**Table 21** 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
					motherboard was replaced. It is also possible that a new server that has the same IP addresses and the same user and password, but was not installed using the AMS flow, exists in the network. EMC does not support such field replacement cases. Such cases should be solved by replacing nodes using EMC-recommended procedures. For more information, contact Customer Support.
RFcache pool state is "Not Started"	XTREMCACHE_INVALID_STATE	Node.Node.XTREMCACHE_INVALID_STATE	SIO07.05.0000010	3 (Error)	Check the Read Flash Cache pool state. The cache device might be misconfigured. In this case, remove the cache device from the cache pool and add it back again.
Host certificate is about to expire	HOST_CERTIFICATE_ABOUT_TO_EXPIRE	Node.Node.HOST_CERTIFICATE_ABOUT_TO_EXPIRE	SIO07.05.0000014	3 (Error)	Renew the AMS certificate and then run "renew certificate" process
CMOS battery state is invalid	NODE_INVALID_CMOS_BATTERY	Node.Node.Invalid_Cmos_Battery	SIO07.05.0000015	3 (Error)	Check the node's CMOS battery. It may need to be replaced.

**Table 21** 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 node is not being managed by the current AMS.	NODE_NOT_MANAGED_BY_AMS	Node.Node_Not_Managed_By_Ams	SIO07.05.000016	5 (Critical)	Verify that the node is not being managed by a different AMS.
Failed to connect to the SVM due to authentication failure	NODE_FAILED_TO_CONNECT_TO_SVM_AUTH_FAILED	Node.Node_Failed_To_Connect_To_Svm_Auth_Failed	SIO07.05.000017	5 (Critical)	Change the SVM's root password to be the one set by the AMS
Failed to connect to the host. The SSH key has been changed, and is not trusted by the AMS.	NODE_FAILED_TO_CONNECT_TO_HOST_SSH_KEY_MISMATCH	Node.Node_Failed_To_Connect_To_Host_Ssh_Key_Mismatch	SIO07.05.000018	5 (Critical)	This is a security issue. Remove the node, and add the node again to the AMS.
Failed to connect to the SVM. The SSH key has been changed, and is not trusted by the AMS.	NODE_FAILED_TO_CONNECT_TO_SVM_SSH_KEY_MISMATCH	Node.Node_Failed_To_Connect_To_Svm_Ssh_Key_Mismatch	SIO07.05.000019	5 (Critical)	This is a security issue. Remove the node, and add the node again to the AMS.
Failed to connect to the host because the certificate is not trusted	NODE_FAILED_TO_CONNECT_TO_ESX_CERTIFICATE_NOT_TRUSTED	Node.Node_Failed_To_Connect_To_Esx_Certificate_Not_Trusted	SIO07.05.000020	5 (Critical)	This is a security issue. Remove the node, and add the node again to the AMS.
Failed to connect to the host due to authentication failure	NODE_FAILED_TO_CONNECT_TO_HOST_AUTH_FAILED	Node.Node_Failed_To_Connect_To_	SIO07.05.000021	5 (Critical)	Change the host's root password to be the one set by the AMS

**Table 21** 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
		Host_Auth_Failed			
The MDM has a self-signed certificate that was not replaced by the AMS. Connection to the MDM is not authenticated by the AMS.	AMS_MDM_HAS_SELF_SIGNED_CERTIFICATE	Node.No de.Ams_Mdm_Has_Self_signed_certificate	SIO07.05.000022	5 (Critical)	Disable AMS secure connection, and enable it again in order to automatically sign the MDM certificate. If the previous solution does not work, remove and add the the MDM again in secure mode.
CPU cache is not enabled	SOCKET_CACHE_DISABLED	Cpu_Socket.Socket_Cache.Disabled	SIO08.01.000001	3 (Error)	Enable the CPU cache in the server BIOS.
CPU cache size is not optimal	SOCKET_CACHE_SIZE_NOT_MAX_SIZE	Cpu_Socket.Socket_Cache.SIZE_NOT_MAX_SIZE	SIO08.01.000002	3 (Error)	Check the CPU cache in the server BIOS, and set it to maximum size.
Unable to connect to ESRS Gateway	ESRS_CONNECTIVITY_ERROR	Esrs.Esrs.CONNECTIVITY_ERROR	SIO10.01.000001	5 (Critical)	Check the network's connectivity to the ESRS Gateway.
The system is not registered with an ESRS gateway	ESRS_NOT_REGISTERED	Esrs.Esrs.NOT_REGISTERED	SIO10.01.000004	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	ESRS_REACHED_CAPACITY_LIMIT	Esrs.Esrs.REACHED_CAPACITY	SIO10.01.000005	5 (Critical)	ESRS has a limit of receiving up to 200 alerts per 8 hours. The limit has been

**Table 21** 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
will be sent in the next 8 hours		ACITY_LIMIT			reached, so no ESRS messages will be sent in the following 8 hours.
The automatic log collection directory is full	AUTOMATIC_LOGS_COLLECT_DIRECTORY_ABOVE_HIGH_THRESHOLD	System. Automatic_Logs_REACHED_CAPACITY_LIMIT	SIO12.01.000003	2 (Warning)	Delete some files from the directory: (Linux) /opt/emc/scaleio/gateway/temp/scaleio-auto-collect-logs/ or (Windows) C:\Program Files\EMC\ScaleIO\Gateway\Temp\sacleio-auto-collect-logs\
There is not enough disk space to run automatic log collection	AUTOMATIC_LOGS_COLLECT_NOT_ENOUGH_DISK_SPACE	System. Automatic_Logs_NOT_ENOUGH_DISK_SPACE	SIO12.01.000004	5 (Critical)	Delete some files from your disk
SDC is disconnected from SDS	ONE_SDC_DISCONNECTED_FROM_ONE_SDS	System. SDC.One_Sdc_Disconnected_From_One_Sds	SIO01.07.000001	5 (Critical)	Check the network links between the affected SDS and SDC
SDC is disconnected from the SDS IP address	ONE_SDC_DISCONNECTED_FROM_ONE_SDS_IP	System. SDC.One_Sdc_Disconnected_From_One_Sds_Ip	SIO01.07.000002	2 (Warning)	Check the network links between the affected SDC and SDS IP addresses.
An SDC is disconnected from all SDSs	ONE_SDC_DISCONNECTED_FROM_ALL_SDS	System. SDC.One_Sdc_Disconnected_From_All_Sds	SIO01.07.000003	5 (Critical)	Check the network links between the affected SDC and all SDSs.

**Table 21** 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
All SDCs are disconnected from the SDS	ALL_SDC_DISCONNECTED_FROM_ONE_SDS	System.SDC.All_Sdc_Disconnected_From_One_Sds	SIO01.07.00 00004	5 (Critical)	Check the network links between all SDCs and the affected SDS.
All SDCs are disconnected from the SDS IP address	ALL_SDC_DISCONNECTED_FROM_ONE_SDS_IP	System.SDC.All_Sdc_Disconnected_From_One_Sds_Ip	SIO01.07.00 00005	2 (Warning)	Check the network links between all SDCs and the affected SDS IP address.
All SDCs are disconnected from all SDSs	ALL_SDC_DISCONNECTED_FROM_ALL_SDS	System.SDC.All_Sdc_Disconnected_From_All_Sds	SIO01.07.00 00006	5 (Critical)	Check the network links between all SDCs and SDSs.
Disconnected network links between SDCs and SDSs.	SDC_MULTIPLE_DISCONNECTIONS_FROM_SDS	System.SDC.Sdc_Multiple_Disconnections_From_Sds	SIO01.07.00 00007	5 (Critical)	Check the network links between all SDCs and SDSs.
This SDC has not been approved	SDC_NOT_APPROVED	System.SDC.Sdc_Not_Approved	SIO01.07.00 00008	2 (Warning)	Add the IP address for this SDC to the list of approved IP addresses.
This SDC does not have an approved IP address	SDC_DOES_NOT_HAVE_APPROVED_IPS	System.SDC.Sdc_Does_Not_Have_Approved_Ips	SIO01.07.00 00009	2 (Warning)	Update the list of approved IP addresses for this SDC.
SDC has an IP address that is not approved by the MDM	SDC_HAS_UNAPPROVED_IP	System.SDC.Sdc_Has_Unapproved_Ips	SIO01.07.00 00010	3 (Error)	Add the IP address for this SDC to the list of approved IP addresses.
Some of the the approved SDC IP	SDC_HAS_UNKNOWN_APPROVED_IP	System.SDC.Sdc	SIO01.07.00 00011	2 (Warning)	Right click the relevant SDC,

**Table 21** 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
addresses in the MDM are not recognized by AMS		_Has_Unknown_Approved_Ips			select "Configure Approved IP Addresses", and add the second data link IP address

## 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. Creating a Lockbox
2. Configuring SNMP after deployment

### 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>
```

#### Note

From system version 2.5 and later, the installation process will assign a random passphrase to this property, and it is highly recommended not to configure or use this property, because it could create a security breach.

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.

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
<code>snmp.sampling_frequency</code>	The MDM sampling period. The default is 30.
<code>snmp.resend_frequency</code>	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

```



# CHAPTER 18

## ScaleIO SDC on AIX server

The ScaleIO SDC can be installed on AIX servers.

This section describes items relevant to using AIX servers in the ScaleIO environment.  
Topics include:

- [SAN virtualization layer](#) ..... 310
- [SDC activities and monitoring](#) ..... 311
- [Enable error logging on AIX servers](#) ..... 313
- [Modify MDM IP address and GUID on AIX server](#) ..... 313

## 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 =  $5 * (\text{gross capacity} - \text{used capacity})$   
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

- When a volume is defined on an AIX SDC, one device is created with the following pathnames:
  - A block device, named /dev/scinidX...n, where X is a number, starting from "0."
  - A raw device, named /dev/rscinidX...n, where X is a number, starting from "0."

In general, mapping SDCs to AIX raw devices will yield best performance. If you are using the device to create a filesystem, use the block device.

- 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.

## SDC activities and monitoring

**Table 22** 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	<code>query_all_sdc</code>	Frontend > SDCs > select the SDC and display the Property Sheet	Used for viewing	
Remove SDC		<code>remove_sdc</code>	Frontend > SDCs > right-click the SDC and select Remove		
Rename an SDC		<code>rename_sdc</code>	Frontend > SDCs > right-click the SDC and select Rename		
Add an MDM definition to an SDC (usually to add the SDC to an additional SIO system)		<code>drv_cfg --add_mdm</code>			The <code>drv_cfg</code> 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
Load an MDM configuration file to an SDC		<code>drv_cfg --load_cfg_file</code>			

**Table 22** SDC activities and monitoring (continued)

To do this	To view this	Use this CLI command	Use the GUI	Use the plug-in	Notes
Modify an existing MDM IP's address configured on an SDC		<code>drv_cfg --mod_mdm_ip</code>			the user documentation for the correct usage.
Modify MDM IP addresses or GUID configured on an ESX-based SDC		<code>esxcli system module parameters list -m scini</code>			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.  <b>Note</b> GUIDs are automatically configured, and modifying them is recommended only for troubleshooting purposes.
	Query GUID and MDM IP addresses on an ESX-based SDC	<code>esxcli system module parameters list -m scini</code>			
	Query SDC state on an ESX-based SDC	<code>esxcli system module list  grep scini</code>			
Abort SDC removal		<code>abort_remove_sdc</code>			
	SDC query	<code>query_sdc</code>			
	Query all active tgt objects	<code>drv_cfg --query_tgts</code>			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	<code>drv_cfg --query_guid</code>			
	Query to determine to which MDM an SDC is connected	<code>drv_cfg --query_mdms</code>			



## Enable error logging on AIX servers

Enable error logging in AIX servers to enhance troubleshooting.

ScaleIO reports errors, diagnostic messages, and failover recovery messages through the syslog file specified by the administrator.

To log messages in `/var/log/messages`:

### Procedure

1. Add the following entry to the `/etc/syslog.conf` file:

```
*.debug /var/log/messages
```

You can also configure `/etc/syslog.conf` to rotate the log file. For example, to rotate the file daily, keep one week's worth of files, and compress files to save space, add the following entry:

```
*.debug /var/log/messages rotate files 7 time 1d compress
```

2. Create the `/var/log/messages` file:

```
touch /var/log/messages
```

3. Enable logging:

```
refresh -s syslogd
```

### Results

Logging is enabled.

## Modify MDM IP address and GUID on AIX server

When SDC is running on an AIX server, you can modify the MDM IP addresses and GUIDs.

GUIDs are assigned automatically, and modifying them should be done with utmost caution.

### Procedure

1. Modify the MDM IP address or GUID.

Option	Description
<b>Persistent</b>	Run this command: <pre>chdev -l scini -a "mdm_ip[1..4]=&lt;&lt;MDM_CLUSTER_IPS&gt;&gt;"</pre>

Option	Description
	<p>Where <i>mdm_ip [1..4]</i> is the cluster number to add or change and <i>MDM_CLUSTER_IPS</i> is a comma-separated list.</p> <p>Examples:</p> <p>To add a new MDM cluster to a newly-installed SDC, with no previous MDM IP addresses assigned:</p> <pre>chdev -l scini -a "mdm_ip1=1.1.1.1,2.2.2.2"</pre> <p>To add an additional MDM cluster to an SDC:</p> <pre>chdev -l scini -a "mdm_ip2=3.3.3.3,4.4.4.4"</pre> <p>To change the MDM IP addresses of a previously-added cluster:</p> <pre>chdev -l scini -a "mdm_ip1=5.5.5.5,6.6.6.6"</pre> <p>To remove the second cluster:</p> <pre>chdev -l scini -a "mdm_ip2="</pre>

## Results

# 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

<b>Degraded Capacity</b>	The capacity is available, but is not protected in case of another failure
<b>Device</b>	Physical storage device, such as a flash drive, or magnetic disk
<b>DirectPath</b>	In ScaleIO documentation, we use the term DirectPath to refer to the VMware vSphere VMDirectPath I/O feature.
<b>DRL</b>	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

<b>Failed Capacity</b>	The capacity is inaccessible due to a failure, and data integrity is at risk
<b>Fault Sets</b>	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.
<b>Forward Rebuild</b>	Data in storage will be rebuilt on another server, due to planned or unplanned shutdown of a server.

## I

<b>ID</b>	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.
<b>IP Role</b>	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

<b>LDAP</b>	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.
<b>Lockbox</b>	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

<b>Management IPs</b>	The IP addresses of the MDMs defined in the system that can be used to access the MDM from CLI, GUI and REST.
<b>Management Port</b>	The Port number used by the MDM for purposes of communicating with the nodes in the ScaleIO network.
<b>Manager MDM</b>	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.
<b>Master MDM</b>	The MDM in the cluster that controls the SDSs and SDCs.
<b>MDM</b>	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

<b>Page Size</b>	The page size, typically in KB, used for caching purposes by Read Flash Cache.
<b>Pass-Through Mode</b>	Data is passed through to or from storage devices without being cached by Read Flash Cache.
<b>Pending Backward Rebuild</b>	A backward rebuild is waiting in a queue, and will be performed when possible, according to rebuild throttling policy.
<b>Primary MDM</b>	See <a href="#">Master MDM</a> .
<b>Protected Capacity</b>	Capacity that has an accessible copy in the system, in case of failure.
<b>Protection Domain</b>	A unique set of SDSs grouped together for reliability and tenancy separation.

## R

<b>RAM Read Cache (RMcache)</b>	Server RAM that is reserved for caching storage devices in a Storage Pool.
<b>Read Flash Cache (RFcache)</b>	Read-only caching of storage devices performed by one or more designated SSD devices and PCIe flash devices in a ScaleIO system.
<b>Rebalance</b>	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.
<b>Rebuild</b>	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.
<b>Restricted MDM Mode</b>	A mode set in which commands can only be performed from an MDM machine.
<b>Restricted SDC Mode</b>	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

