

Dell EMC ScaleIO

Version 2.x

Deployment Guide

P/N 302-003-287

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

CONTENTS

Figures	9
Tables	11
Preface	13
Part 1 Introduction	15
Chapter 1 Introduction to EMC ScaleIO	17
What is ScaleIO?	18
System requirements	18
ScaleIO cluster components	19
Physical server requirements	19
Supported operating systems	20
External SDC support	22
GUI server requirements	22
ScaleIO Gateway server requirements	23
Other requirements	23
New and changed features	24
New features for ScaleIO	24
Changed features	24
Product limits	24
Chapter 2 Architecture	27
ScaleIO Architecture Overview	28
System	28
Hardware	28
Software	28
The MDM cluster	29
Storage definitions	31
Protection Domains	31
Storage Pools	31
Fault Sets	33
Naming	34
Protection and load balancing	34
Rebuild	34
Rebuild throttling	35
Rebalance	35
Rebalance throttling	36
Checksum protection	37
ESX vStorage APIs for Array Integration (VAAI)	37
Caching	38
Networking	41
VMware limitation	44
Virtual IP Address	44
Monitoring of SDC and SDS connections	45
S.M.A.R.T. hardware monitoring	46

	List of approved RAID controllers.....	47
	Snapshots.....	48
	V-Trees.....	49
	Other functions.....	49
	Implementing ScaleIO.....	53
	Physical layer.....	53
	SAN virtualization layer.....	55
	Implementing ScaleIO over a virtual system.....	56
	Implementing ScaleIO in an ESXi-based system.....	56
	Xen implementation.....	59
	Maintenance.....	59
	Maintaining the physical layer.....	59
	Instant maintenance mode.....	60
	Maintaining the virtualization layer.....	60
	Management tools.....	61
	Configuring direct attached storage (DAS).....	61
Part 2	Deploying ScaleIO Systems	63
Chapter 3	Deploying ScaleIO on Physical Servers	65
	Overview of ScaleIO deployment on physical servers.....	66
	Which deployment mode is best for you?.....	67
	Linux package names.....	67
	Preparing the Installation Manager and the Gateway.....	68
	Preparing the IM on a Linux server.....	69
	Preparing the IM on a Windows server.....	71
	Installing with the full Installation Manager.....	72
	Preparing the CSV topology file.....	72
	Installing with the Installation Manager.....	77
	Installing with the Installation Manager wizard.....	84
	Setting up the installation.....	85
	Running the installation.....	88
	Enabling the storage.....	89
	Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers.....	91
	Ensuring the kernel version is correct.....	92
	Creating the configuration file via the ScaleIO Gateway.....	93
	Creating the configuration file manually.....	95
	Update the ScaleIO signature key.....	96
	Running the ScaleIO Gateway on CoreOS container.....	97
	Creating a mirror repository.....	97
	Deploying on OpenStack.....	98
	ScaleIO with Mirantis OpenStack.....	98
	ScaleIO with Canonical (Ubuntu) OpenStack.....	99
	ScaleIO with OpenStack Liberty.....	99
	ScaleIO with OpenStack Mitaka.....	100
	Environment and system requirements.....	101
	Install the ScaleIO GUI.....	102
	Configuring Installation Manager properties.....	102
Chapter 4	Deploying ScaleIO on ESX Servers	105
	Automatic Deployment of ScaleIO on ESXi servers.....	106
	Deployment prerequisites.....	106
	Prepare the ESXi environment.....	109
	Registering the ScaleIO plug-in.....	109

	Uploading the OVA template.....	111
	Accessing the plug-in.....	112
	Preparing the ESXi hosts.....	113
	Deploying ScaleIO with DirectPath device management.....	114
	Adding devices to an SDS.....	123
	Deploying ScaleIO with RDM/VMDK device management.....	124
	Install the ScaleIO GUI.....	134
	Rolling back the deployment wizard in the vSphere Web plug-in.....	135
Chapter 5	Post-Deployment Activities	137
	Post-deployment checklist.....	138
	Create a Lockbox.....	139
	Configuring ESRS.....	140
	Configure native users.....	141
	Configure LDAP users.....	142
	Configuring SNMP after deployment.....	143
	Create and map volumes.....	144
	Creating and mapping volumes using the CLI.....	144
	Creating and mapping volumes using the GUI.....	144
	Creating and mapping volumes using the vSphere plug-in.....	145
Chapter 6	Licensing	147
	Licensing overview.....	148
	Activating entitlements and installing a license file.....	149
	Activating an entitlement and downloading the license file.....	149
	Installing the license.....	153
	License file example.....	153
	Error messages.....	154
Part 3	Reference	155
Chapter 7	Manual Deployment	157
	Manual deployment on physical servers.....	158
	MDM cluster setup.....	160
	Post-installation considerations.....	172
	Provisioning storage example.....	177
	Manual Deployment of ScaleIO on ESXi Servers.....	180
	Deploying the ScaleIO virtual machine (SVM).....	180
	Configuring the UUID on virtual machines.....	182
	Installing the SDC directly on an ESX host.....	182
	Optimizing the guest operating system performance.....	183
	Manually removing the VIB file from the ESX host.....	184
Chapter 8	Advanced Topics	187
	Advanced Gateway topics.....	188
	Enabling and disabling Gateway components.....	188
	Using a custom Java configuration for the Gateway.....	189
	Installing the Gateway without assigning an admin password.....	189
	Certificate management for ScaleIO Gateway.....	189
	OpenStack interoperation with the ScaleIO Gateway.....	193
	Generating a self-signed certificate using the keytool utility.....	193
	Advanced Installation Manager topics.....	193

	Adding devices to SDS nodes on Windows servers.....	193
	Installing without validating Linux devices.....	195
	Using SSH authentication on the ScaleIO Gateway.....	195
	Using the IM REST API.....	196
	Advanced vSphere plug-in topics.....	213
	Advanced settings options.....	213
	Registering the ScaleIO plug-in manually.....	214
	Troubleshooting plug-in registration issues.....	216
	Deploy ScaleIO in a 2-layer environment.....	217
	Deploying ScaleIO in a 2-layer environment using the Installation Manager and plug-in.....	217
	Manually deploying ScaleIO in a 2-layer environment.....	218
	Prepare the servers for deployment.....	219
Chapter 9	Maintaining a ScaleIO System	223
	Extending the MDM cluster from 3 to 5-node.....	224
	Extending the MDM cluster in physical servers.....	224
	Extending the MDM cluster in VMware servers.....	225
	Extending the MDM cluster.....	225
	Updating the SDC parameters.....	226
	Installing RFcache on servers in an existing ScaleIO system.....	227
	Installing RFcache in physical servers.....	227
	Enabling RFcache on VMware servers.....	227
	Creating a Lockbox for SNMP, ESRS, or LDAP.....	228
	Switching to secured authentication mode.....	228
	Physical Linux servers.....	228
	Physical Windows servers.....	229
	VMware servers.....	230
	Working with Dynamic Host Name resolution for SNMP in ScaleIO.....	231
	Using SCLI in non-secure mode.....	233
	Extending an existing ScaleIO system.....	233
	Adding components with the Installation Manager.....	234
	Adding components with the VMware deployment wizard.....	234
	Configuring virtual IP addresses using Installation Manager.....	235
	Removing ScaleIO.....	236
	Removing ScaleIO using the IM.....	236
	Unregistering the ScaleIO plug-in.....	237
Chapter 10	System Analysis	239
	System analysis overview.....	240
	Creating the system analysis report.....	241
	System analysis report description.....	243
Chapter 11	Configuring ESRS connection properties	249
	Before configuring ESRS.....	250
	Registering the system with the ESRS Gateway.....	250
	Performing other ESRS configuration activities.....	252
	Validating ConnectEMC Dial-Home.....	253
Chapter 12	Common Tasks	255
	Install the ScaleIO GUI.....	256
	Log in to the ScaleIO GUI.....	256
	Connection and disconnection information.....	257

Add LIA to a system to enable automated upgrade.....	257
Associating ScaleIO volumes with physical disks.....	258
Volume information - Linux.....	258
Volume information - Windows.....	259
Volume information - AIX.....	260
Port usage and changing default ports.....	261
Adding an external SDC to an existing ScaleIO system.....	262
Installing SDC on an ESX server and connecting it to ScaleIO.....	262
Installing SDC on a Linux server and connecting it to ScaleIO.....	263
Install SDC on an AIX server and connect it to ScaleIO.....	264
Installing SDC on a Windows server and connecting it to ScaleIO....	264
Changing the LIA configuration file.....	265
Cleaning the ScaleIO VMware environment and performing a clean install....	266
Configuring ScaleIO devices in Linux LVM.....	267
Configuring session timeout parameters.....	268
Fixing keytool errors.....	268
Error during rpm installation command.....	268
Error during rpm upgrade command.....	269
Installing Java on SUSE 12 servers.....	269
SVM manual memory allocation.....	269
Upgrading Java.....	271
Mounting ScaleIO.....	271
The ScaleIO Gateway web server isn't responding.....	273
The ScaleIO Gateway (REST service, Installation Manager) may be disabled:.....	273
The ScaleIO Gateway web server isn't responsive and the following error appears in the catalina log file:.....	273
Upgrading the Gateway when a custom certificate is used.....	274
Uploading a new OVA.....	275
Using the same data network for different NICs.....	275
What to do when the default self-signed certificate expires.....	275
Add another IP address subnet to an MDM cluster.....	275
Shutdown or restart a node gracefully.....	277
Gracefully shut down or reboot a node.....	277
Return the node to operation.....	280
Deployment of ScaleIO using a non-root user.....	280
Configure a non-root non-interactive sudo user.....	281

Glossary

283

FIGURES

1	5-node MDM cluster.....	30
2	Protection Domains and Storage Pools.....	32
3	Protection Domains, Storage Pools, and Fault Sets.....	33
4	ScaleIO system deployed on a single network.....	43
5	ScaleIO system deployed on separate networks.....	44
6	Snapshot operations.....	48
7	V-Tree diagram.....	49
8	Physical layout example—3-node cluster.....	53
9	Physical layout example—3-node cluster.....	54
10	ScaleIO implementation on ESX.....	57
11	ScaleIO Xen virtual machine architecture.....	59
12	CSV—complete.....	73
13	OpenStack Liberty integration.....	100
14	OpenStack Mitaka integration.....	101
15	EMC ScaleIO screen.....	112
16	Licensing LAC email.....	149
17	License file example.....	153
18	Example of IM REST API URI when JSESSION ID is configured as a header.....	212
19	Before extending.....	224
20	After extending.....	225
21	Set Virtual IPs for ScaleIO system screen.....	235

FIGURES

TABLES

1	Server physical requirements.....	19
2	Supported operating systems - ScaleIO components.....	21
3	Product limits.....	24
4	MDM cluster modes.....	30
5	Caching modes.....	38
6	Caching support matrix.....	40
7	IP address configurations in ScaleIO (based on CSV file).....	42
8	Linux package formats.....	67
9	CSV topology spreadsheets.....	73
10	driver_sync.conf parameters.....	95
11	eLicensing terminology(continued).....	148
12	Licensing error messages.....	154
13	Sample manual installation topology.....	160
14	Sample manual installation topology.....	166
15	Provisioning storage example.....	178
16	IP address table for manual deployment on ESXi servers.....	180
17	Response.....	198
18	Default ports.....	261

Preface

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

Contact your EMC technical support professional if a product does not function properly or does not function as described in this document.

Note

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

Related documentation

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

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

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

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

Typographical conventions

EMC uses the following type style conventions in this document:

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

PART 1

Introduction

This part gives an overview of ScaleIO benefits and architecture. Chapters include:

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

[Chapter 2, "Architecture"](#)

CHAPTER 1

Introduction to EMC ScaleIO

This chapter introduces EMC® ScaleIO®.

This guide is specific to ScaleIO software deployments.

For ScaleIO Ready Node or VxRack Node 100 Series systems, use this guide as a general reference, and refer to your product's Hardware Configuration and Operating System Installation Guide and the Quick Start Guide for step-by-step installation direction.

Topics include:

• What is ScaleIO?	18
• System requirements	18
• New and changed features	24
• Product limits	24

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 RfCache will be enabled.
- ScaleIO Gateway server on a separate server, outside of the ScaleIO system.

Physical server requirements

Table 1 Server physical requirements

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

Table 1 Server physical requirements (continued)

Component	Requirement
	<ul style="list-style-type: none"> 2U1N servers—1 GB RAM <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"> 1 GB for each physical node or Xen hypervisor 10 GB for VMware topologies
Connectivity	<p>One of the following:</p> <ul style="list-style-type: none"> 1 GbE or 10 GbE (recommended) network IP-over-InfiniBand network <p>Dual-port network interface cards (recommended)</p> <p>Ensure the following:</p> <ul style="list-style-type: none"> There is network connectivity between all components. Network bandwidth and latency between all nodes is acceptable, according to application demands. Ethernet switch supports the bandwidth between network nodes. MTU settings are consistent across all servers and switches. The following TCP ports are not used by any other application, and are open in the local firewall of the server: <ul style="list-style-type: none"> MDM: 6611 and 9011 Tie Breaker: 9011 SDS: 7072. Multiple SDS (not supported on Windows): 7073-7076 Light Installation Agent (LIA): 9099 SDBG ports (used by ScaleIO internal debugging tools to extract live information from the system): MDM 25620, SDS 25640. Multiple SDS (not supported on Windows): 25641-25644 (not 25640). The following UDP port is open in the local firewall of the server: <ul style="list-style-type: none"> SNMP traps: 162 <hr/> <p>Note</p> <p>You can change the default ports. For more information, see “Changing default ports” in the user documentation.</p>

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/elc/elcnavhome>.

Table 2 Supported operating systems - ScaleIO components

Operating system	Requirement
Linux	<p>Supported versions:</p> <ul style="list-style-type: none"> CentOS 6.x-7.x, Oracle Linux 6.5/7.x Red Hat 6.x-7.x SUSE 11.3, 11.4, 12, 12.1, 12.2 Ubuntu 14.04, Ubuntu 16.04 <hr/> <p>Note</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"> numactl libaio <p>Additional packages required for MDM components:</p> <ul style="list-style-type: none"> bash-completion (for SCLI completion) Latest version of Python 2.X <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"> <code>libopenssl1_0_0-1.0.1g-0.40.1.x86_64.rpm</code> <code>openssl1-1.0.1g-0.40.1.x86_64.rpm</code> <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"> 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.) For VxRack Node 100 Series, only 2012 R2 is supported. 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 https://support.emc.com. To install SDC or RFCache on 2008 R2, ensure that Microsoft Security Update KB3033929 is installed. <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"> OpenSSL 64-bit v1.0.1 or later (v1.1, however, is not supported)

Table 2 Supported operating systems - ScaleIO components (continued)

Operating system	Requirement
	<ul style="list-style-type: none"> Visual C++ redistributable 2010 package, 64-bit <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"> VMware ESXi OS: 5.5 U3, 6.0 U3, or 6.5, managed by vCenter 5.5, 6.0, or 6.5 Hyper-V XenServer 6.5 or 7.0 <hr/> <p>Note</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"> Red Hat KVM

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"> UNIX: AIX 7.2 hLinux: 4.x/5.x

GUI server requirements

Component	Requirement
Supported operating systems	<ul style="list-style-type: none"> Windows: <ul style="list-style-type: none"> 7, 2008 R2, 10, 2012 or 2012 R2, 2016. Server Core editions are not supported. Linux: <ul style="list-style-type: none"> CentOS 6.x-7.x, Oracle Linux 6.5/7.x Red Hat 6.x-7.x SUSE 11.3, 11.4, 12, 12.1, 12.2 Ubuntu 14.04, Ubuntu 16.04
Other	<ul style="list-style-type: none"> v1.8 (64-bit), build 149 or earlier. <p>You can download previous versions from this link: http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html</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"> Screen resolution: 1366 x 768 minimum

ScaleIO Gateway server requirements

Component	Requirement
Supported operating systems	<ul style="list-style-type: none"> Windows 2008 R2, 2012 R2, or 2016, including the Visual C++ redistributable 2010 package, 64-bit. Server Core editions are not supported. Linux: <ul style="list-style-type: none"> CentOS 6.x-7.x Oracle Linux 6.5/7.x Red Hat 6.x-7.x SUSE 11.3, 12, 12.1, and 12.2 Ubuntu 14.04, Ubuntu 16.04 <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"> Internet Explorer 10, or later Firefox, version 42, or later Chrome, version 45, or later
Java	<ul style="list-style-type: none"> v1.8 (64-bit), build 149 or earlier. <p>You can download previous versions from this link: http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html</p>
Other	<ul style="list-style-type: none"> For a Windows Gateway, the Windows Management Instrumentation service must be enabled on the IM server and on all Windows ScaleIO nodes. Do not install the Gateway on a server on which RFCache will be enabled or on which SDC will be installed. The Gateway server must have connectivity to all the nodes that are being installed. If you are using separate networks for management and data, the server must be able to communicate with both networks.

Other requirements

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

NOTICE

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

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

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

a. If more are needed, contact EMC Support.

b. On VMware servers, the maximum devices per SDS is 59.

c. Replication support is version-specific. For information, see the ESSM.

CHAPTER 2

Architecture

This chapter describes the ScaleIO architecture. Topics include:

• ScaleIO Architecture Overview	28
• System	28
• The MDM cluster	29
• Storage definitions	31
• Protection and load balancing	34
• ESX vStorage APIs for Array Integration (VAAI)	37
• Caching	38
• Networking	41
• Virtual IP Address	44
• Monitoring of SDC and SDS connections	45
• S.M.A.R.T. hardware monitoring	46
• Snapshots	48
• Other functions	49
• Implementing ScaleIO	53
• Implementing ScaleIO over a virtual system	56
• Maintenance	59
• Management tools	61
• Configuring direct attached storage (DAS)	61

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

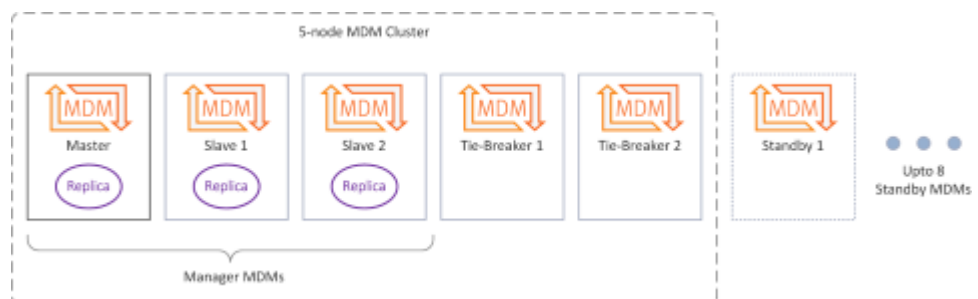
Cluster mode	Members	Description
3-node (default)	<ul style="list-style-type: none"> • Master MDM • Slave MDM • Tie Breaker 	3-node cluster has two copies of the repository, thus can withstand one MDM cluster failure.
5-node	<ul style="list-style-type: none"> • Master MDM • Two Slave MDM • Two Tie Breaker 	5-node cluster has three copies of the repository, thus can withstand two MDM cluster failure.
Single-node	<ul style="list-style-type: none"> • Master MDM 	Single-node cluster has only one copy of the repository, thus it cannot withstand failure. It is not recommended to use Single Mode in production systems, except in temporary situations.

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

The 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:

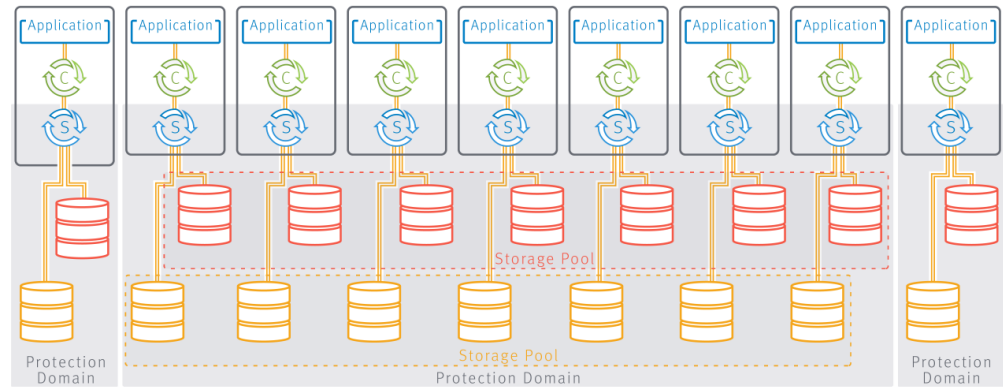
- optimal performance
- reduction of theoretical *mean time between failure* issues
- ability to sustain multiple failures in different Protection Domains

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).

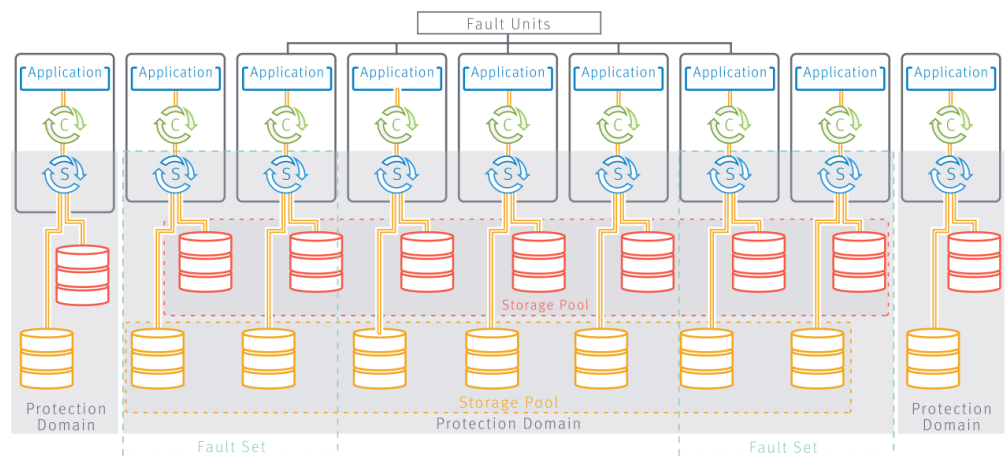
There must be enough capacity within at least 3 fault units to enable mirroring.

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 33 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 RMcach (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>Note</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"> • An RFcache device (flash device) can be partitioned only on Linux. • An SDS storage/source device cannot be partitioned if it needs to be accelerated by RFcache. • An SDS storage/source device as a file (over file system), cannot be accelerated by RFcache. 	
CacheCade	Read and write-back caching performed by one or more dedicated SSD devices in the server.	<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>Note</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>Note</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.

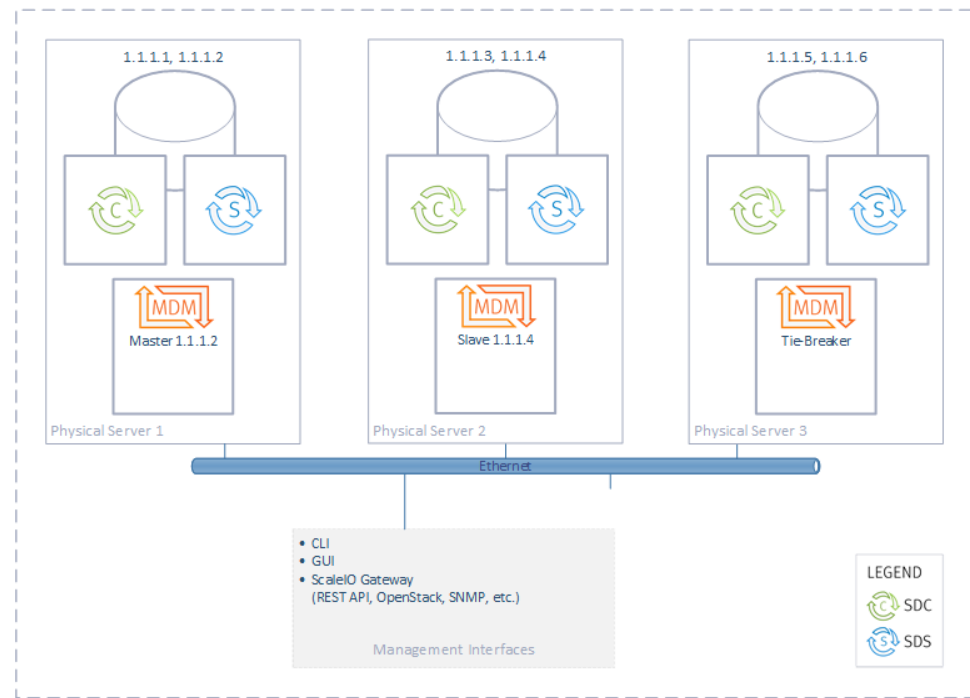
On VMware, all ScaleIO VMs need to have a management IP address as well as another address for the data network, the network on which traffic flows between SDSs and SDCs for read/writes, rebuild, and rebalance.

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:

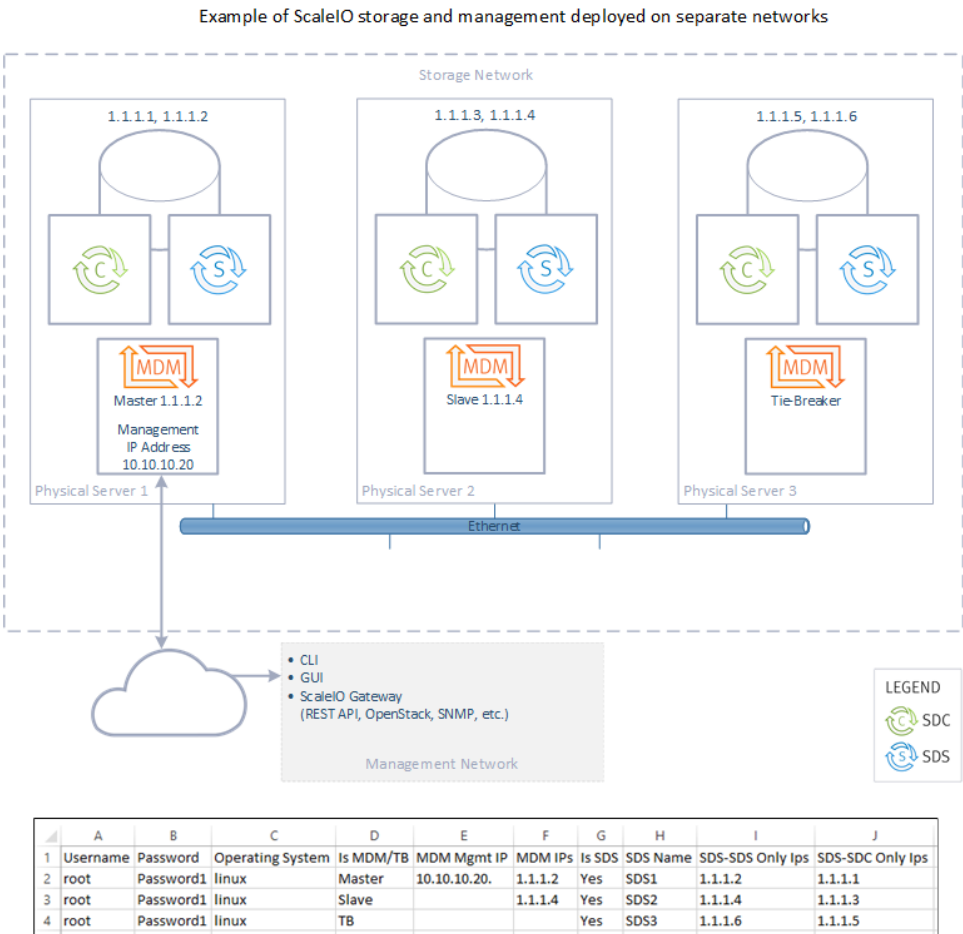
Figure 4 ScaleIO system deployed on a single network

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

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 User Guide*.

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

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

List of approved RAID controllers

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

ScaleIO-certified RAID controllers

The following table describes the ScaleIO-certified RAID controllers:

Manufacturer	Specifications
HP	<ul style="list-style-type: none"> Model Name: Smart Array P440ar Vendor Name: HP Firmware Version: 3.56 Driver Version: 3.4.10 Driver Name: hpsa PCI Address: 0000:03:00.0
DELL	<ul style="list-style-type: none"> Model Name: PERC H730 Mini Vendor Name: Dell Firmware Version: 25.3.0.0016 Driver Version: 06.807.10.00-rh Driver Name: megaraid_sas PCI Address: 00:02:00:00
LSI	<ul style="list-style-type: none"> Model Name: LSI MegaRAID SAS 9271-8i Vendor Name: LSI Firmware Version: 23.12.0-0021 Driver Version: 06.810.09.00-rh Driver Name: megaraid_sas PCI Address: 00:82:00:00
	<ul style="list-style-type: none"> Model Name: LSI MegaRAID SAS 9271-8i Vendor Name: LSI Firmware Version: 23.12.0-0018 Driver Version: 06.805.06.01-rc Driver Name: megaraid_sas PCI Address: 00:82:00:00
	<ul style="list-style-type: none"> Model Name: LSI MegaRAID SAS 9271-8i Vendor Name: LSI Firmware Version: 23.12.0-0021 Driver Version: 06.805.06.01-rc Driver Name: megaraid_sas PCI Address: 00:82:00:00

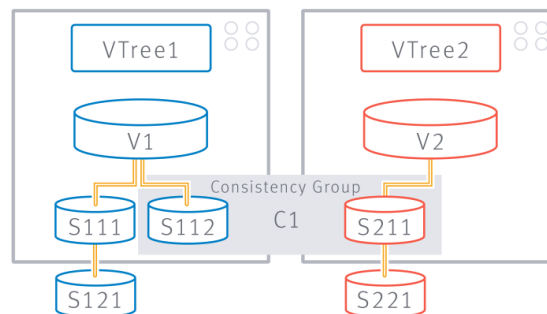
Manufacturer	Specifications
	<ul style="list-style-type: none"> Model Name: LSI MegaRAID SAS 9271-8i Vendor Name: LSI Firmware Version: 23.34.0-0005 Driver Version: 06.810.08.00 Driver Name: megaraid_sas PCI Address: 00:82:00:00

Snapshots

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

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

Figure 6 Snapshot operations



The structure related to all the snapshots resulting from one volume is referred to as a 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 48, in RED, S211 is a snapshot of V2. Since S112 and S211 were taken together, they compose a consistency group designated as C1.

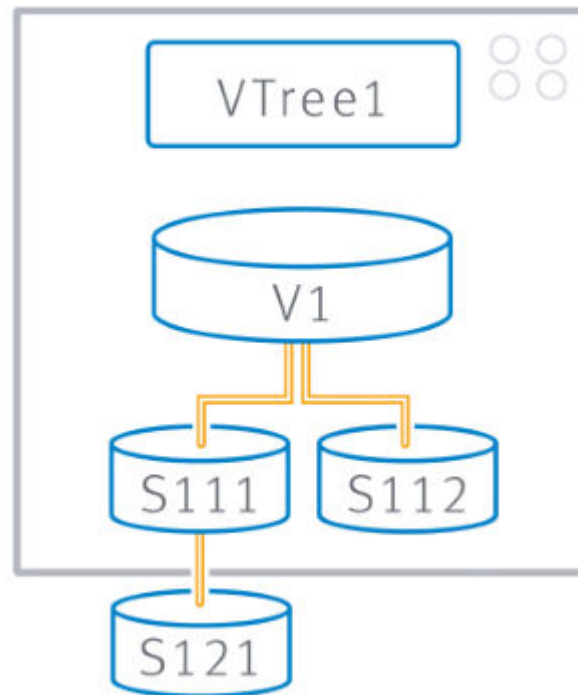
Note

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

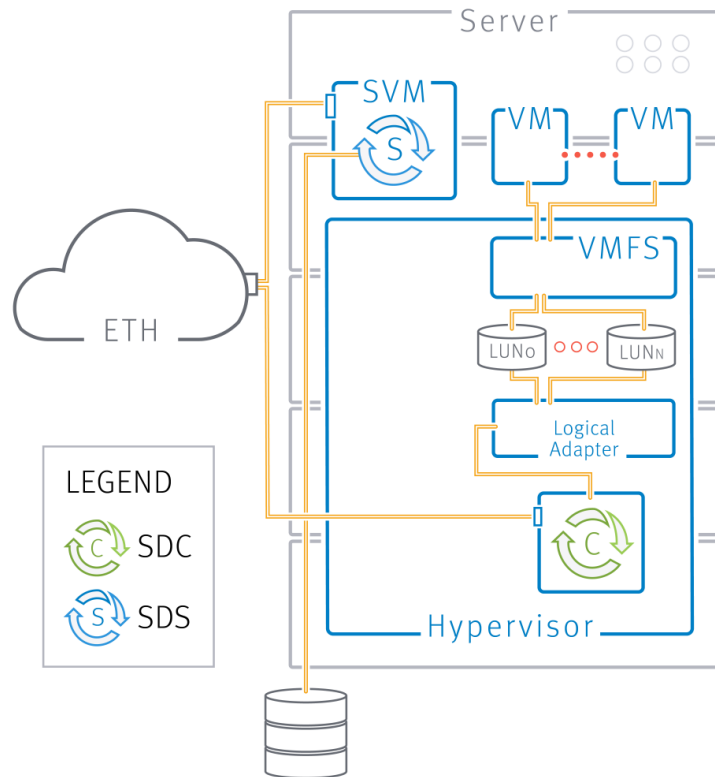
In the VMware environment, the MDM and SDS components are installed on a dedicated SVM, whereas the SDC is installed directly on the ESX host.

Note

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

This implementation is illustrated in the following figure:

Figure 10 ScaleIO implementation on ESX



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

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

Device management

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

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

Requirements:

- When DirectPath is configured, all servers in the system must use DirectPath. For ScaleIO Ready Node or VxRack Node 100 Series systems, this can be configured with the ScaleIO vSphere plug-in. For other servers, use the vSphere client to configure each ESX host manually.
- If the host has multiple controllers, you must configure DirectPath on that host manually, not with the plug-in. After that is done, you can use the plug-in to deploy ScaleIO on that host and to add devices to it.
- All devices on all servers must not have any other use (not VMDK, RDM, or be part of a datastore).
- PowerEdge 14G servers with any NVME devices cannot be added in a DirectPath-based system. Use RDM-based, instead.

- ESX boot device requirements:
 - Must be on a separate controller or connected directly to the board.
 - SATADOM and M2 boot devices are supported. These devices allow the creation of a datastore on the system disk, which is needed to host the ScaleIO VM.
 - USB boot devices are not supported.

- RDM

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

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

Before enabling this feature, contact EMC Support.

- VMDK

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

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

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

Note

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

System size

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

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

Pre-deployment considerations

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

- Do you want to use separate networks for data and management (recommended), and which IP addresses will you use for the SVMs and VMkernels?
- Are there flash devices that will be added to the SDS?
- Do you want to create Fault Sets? See the requirements in [Fault Sets](#) on page 33.

Post-deployment considerations

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

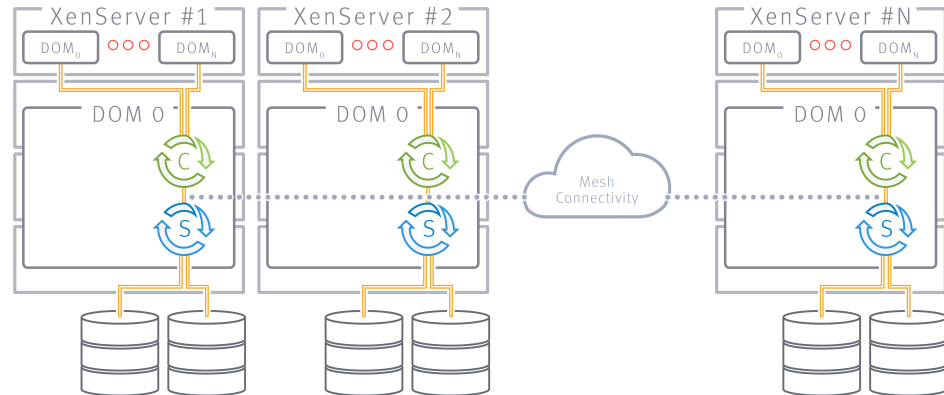
- After deployment is complete, set all SVMs to start automatically with the system. Do not set SVMs under the VMware resource-pool feature.

- 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 interoperability between a ScaleIO system and an OpenStack cloud operating system.
- **REST Gateway**
A REST API can be used to expose monitoring and provisioning via the REST interface. The REST server is installed as part of the ScaleIO Gateway.
Many ScaleIO activities can be performed in more than one management tool.
The following tool is also provided:
- **Installation Manager (IM)**
The IM is used for installing ScaleIO, upgrading and uninstalling components, as well as running the get-info operation. The IM is installed as part of the ScaleIO Gateway.

Configuring direct attached storage (DAS)

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

NOTICE

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

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

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

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

Note

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

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

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

PART 2

Deploying ScaleIO Systems

This part describes how to deploy ScaleIO in your environment. Chapters include:

[Chapter 3, "Deploying ScaleIO on Physical Servers"](#)

[Chapter 4, "Deploying ScaleIO on ESX Servers"](#)

[Chapter 5, "Post-Deployment Activities"](#)

[Chapter 6, "Licensing"](#)

CHAPTER 3

Deploying ScaleIO on Physical Servers

This chapter describes how to install, deploy, and perform initial configuration of ScaleIO software components on physical servers. Topics include:

• Overview of ScaleIO deployment on physical servers	66
• Which deployment mode is best for you?	67
• Linux package names	67
• Preparing the Installation Manager and the Gateway	68
• Installing with the full Installation Manager	72
• Installing with the Installation Manager wizard	84
• Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers	91
• Deploying on OpenStack	98
• Install the ScaleIO GUI	102
• Configuring Installation Manager properties	102

Overview of ScaleIO deployment on physical servers

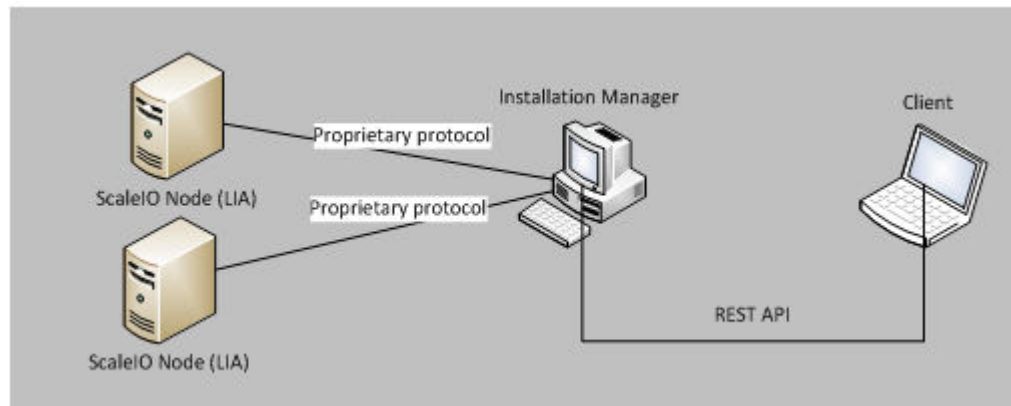
This section describes an overview of ScaleIO deployment on physical servers.

Note

The 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 EMC Support at <https://support.emc.com>.

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

In physical environments, and in non-VMware virtual environments (such as Xen and KVM), you use the ScaleIO Installation Manager (IM) to install and configure the ScaleIO components on multiple nodes from one central server, via a web client.



The Installation Manager can be used to install numerous ScaleIO systems, from one central workstation.

Note

The IM can be disabled and configured, as described in [“Configuring the Installation Manager”](#).

As part of the installation process, the ScaleIO Lightweight Installation Agent (LIA) is installed. The LIA is a required component for many maintenance operations. During initial installation, the IM establishes trust with the LIA, which facilitates its operation.

Note

To make changes to the LIA configuration file, see [“Changing the LIA configuration file”](#).

You can limit the operations that LIA can perform by making changes to its configuration file.

In Windows environments, LIA will only install and upgrade packages that are signed by ScaleIO and associated with the ScaleIO system.

The Installation Manager has a REST API that enables install, extend, and uninstall functionality. These features are described in the *EMC ScaleIO User Guide*.

Which deployment mode is best for you?

You can deploy ScaleIO on physical servers using the Installation Manager or manually.

For manual deployment, see the "Manual Deployment" chapter. The Installation Manager (IM) can be used in the following modes:

- **Full IM**
This mode, which enables the highest level of customization, uses a user-prepared CSV topology file to install ScaleIO and configure nodes. The CSV file is used by the IM to set up and configure all the nodes. To add additional servers after the initial installation, you will also use the combination of CSV topology file and IM.
- **IM wizard**
This abbreviated mode enables you to get a ScaleIO system up and running in the simplest manner, with preset node configuration, where all management and data communication are on the same network. This mode is perfect for a single Protection Domain, fully-converged system. No CSV file is required.

After preparing the IM, you can choose which mode to use.

To prepare the ScaleIO Installation Manager (IM), proceed to [Preparing the Installation Manager and the Gateway](#) on page 68.

Linux package names

ScaleIO Linux installation packages have several name formats.

Package formats

Throughout this section, Linux packages are displayed in the following format:

EMC-ScaleIO-*<component>*-2.5-*<build>*.X.*<flavor>*.x86_64.rpm

where:

- *<component>* is the ScaleIO component: mdm, lia, sds, etc.
- *<flavor>* is the OS for your environment, according to the following table:

Table 8 Linux package formats

Linux flavor	Package format	Example
CoreOS	CoreOS	EMC-ScaleIO-mdm-2.5- <i><build></i> .X.CoreOS.x86_64.tar CoreOS packages may need to be extracted before use. This is described where relevant.
hLinux (SDC only)	hpelinux	EMC-ScaleIO-sdc-2.5- <i><build></i> .X.hpelinux.amd64.deb The hLinux package needs to be extracted before use. This is described where relevant.
RHEL/OL/CentOS	el <i><version></i>	EMC-ScaleIO-mdm-2.5- <i><build></i> .X.el <i><version></i> .x86_64.rpm Example: EMC-ScaleIO-mdm-2.5- <i><build></i> .X.el6.x86_64.rpm
SUSE	sles <i><version></i>	EMC-ScaleIO-mdm-2.5- <i><build></i> .X.sles <i><version></i> .x86_64.rpm Example: EMC-ScaleIO-mdm-2.5- <i><build></i> .X.sles11.3.x86_64.rpm
Ubuntu	Ubuntu. <i><version></i>	EMC-ScaleIO-mdm-2.5- <i><build></i> .X.Ubuntu. <i><version></i> .x86_64.tar Example: EMC-ScaleIO-mdm-2.5- <i><build></i> .X.Ubuntu.16.04.x86_64.tar

Table 8 Linux package formats (continued)

Linux flavor	Package format	Example
		Ubuntu packages may need to be extracted before use. This is described where relevant.
XenServer	xs<version>	EMC-ScaleIO-mdm-2.5-<build>.X.xs<version>.x86_64.rpm Example: EMC-ScaleIO-mdm-2.5-<build>.X.xs7.0.0.x86_64.rpm

Use the packages and the installation commands that match your Linux operating system environment.

On XenServer servers, the syntax for ScaleIO CLI commands is `siocli`, as opposed to `scli`.

Note

Before deploying or upgrading CoreOS, hLinux, Oracle Linux (OL), or Ubuntu systems, ensure that the ScaleIO environment is prepared, as described in [Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers](#) on page 91.

Preparing the Installation Manager and the Gateway

Before installing ScaleIO, you must prepare the Installation Manager (IM). The Installation Manager is installed as part of the ScaleIO Gateway. The Gateway also includes the REST Gateway and SNMP trap sender functionality. In this document, the terms IM server and Gateway server are synonymous.

You can enable and disable Gateway components, as described in [“Enabling and disabling Gateway components”](#).

The Gateway can be installed on the same node as an MDM or SDS ScaleIO component. Do not install the Gateway on an SDS upon which RfCache will be enabled or on an SDC server. Ensure that the node has adequate memory to run the IM (minimum of 3 GB) and any other applications.

To determine on which node to install the Gateway, determine your operating system environment:

ScaleIO node OS	Install Gateway on this OS
Linux	Linux or Windows
Windows	Windows
Mixed (Linux and Windows)	Windows

The IM server must have connectivity to both the data and management ScaleIO networks.

Ensure that your IM server meets the system requirements described in [System requirements](#) on page 18.

The Gateway cannot be installed on a CoreOS server. In CoreOS environment, deploy the Gateway on a supported Windows or Linux server.

Proceed to the section that matches the operating system of your IM server.

Preparing the IM on a Linux server

Prepare the Installation Manager (IM) on a Linux server.

Before you begin

Before deploying any ScaleIO components on a server that was part of another ScaleIO system, use the operating system tools to remove all of those ScaleIO components.

You can use a Linux IM to deploy to Linux servers only.

Procedure

1. Download (and extract, if necessary) the files for your operating system.

You can download all files from the Complete Software download for your software version, available from Online Support (https://support.emc.com/products/33925_ScaleIO-Software).

The following table lists the files:

Required files	<ul style="list-style-type: none"> • Core files. You need these files for all operating systems of nodes in the system. <ul style="list-style-type: none"> ▪ LIA ▪ MDM ▪ SDC ▪ SDS • Gateway, for the operating system of the Gateway server
Optional files	<ul style="list-style-type: none"> • GUI, for the operating system of the GUI server • XCACHE (for RFcache)

To install with a customized Java configuration, see [“Using a custom Java configuration for the Gateway”](#).

2. From the extracted download file, install the ScaleIO Gateway on the Linux IM server, by running the following command (all on one line):
 - RHEL/Centos/Oracle Linux/SLES

```
GATEWAY_ADMIN_PASSWORD=<new_GW_admin_password> rpm -i /tmp/
EMC-ScaleIO-gateway-2.5-<build>.X.x86_64.rpm (for SLES 11.3
add --nodeps)
```

- Ubuntu

```
GATEWAY_ADMIN_PASSWORD=<new_GW_admin_password> dpkg -i /tmp/
EMC-ScaleIO-gateway-2.5-<build>.X.amd64.deb
```

where *<new_GW_admin_password>* is a password that you define to access the IM.

The password must meet the following criteria:

- At least 8 characters long
- Include at least 3 of the following groups: [a-z], [A-Z], [0-9], special chars (!@#\$...)

To install without assigning a Gateway password, see [“Installing the Gateway without assigning an admin password”](#).

Note

To include special characters in the password, such as ! @ \$, when assigning the password, add the "\" character before the special character. For example, to assign the password Scaleio\$\$, enter it as Scaleio\\\$\$. When logging in, the "\" character is not used.

Note

For installation on Ubuntu, Oracle Linux, or CoreOS servers, read this section before continuing with the deployment: [Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers](#) on page 91.

After you finish

Continue with your choice of IM methods:

- To use the IM wizard, skip to [“Installing with the Installation Manager wizard”](#).
 - To use the full IM to install ScaleIO only, continue with [“Installing with the full Installation Manager”](#).
 - To use the full IM to install ScaleIO and the replication splitter for RecoverPoint, continue with [“Installing the replication splitter for RecoverPoint”](#).
-

Note

Replication support is version-specific; see the EMC Simple Support Matrix (ESSM) at <https://elabnavigator.emc.com>.

Note

You can configure the Installation Manager, as described in [“Configuring the Installation Manager”](#)

Installing the replication splitter for RecoverPoint

Note

If you are not installing the replication splitter for RecoverPoint, skip to [“Installing with the full Installation Manager”](#).

This topic describes steps that are required when using the IM to install the replication splitter for RecoverPoint.

Note

Replication support is version-specific; see the ESSM for full details.

There are additional post-installation configurations steps. For full instructions, refer to the ScaleIO Write Splitter for RecoverPoint Technical Notes before beginning the installation.

To install the replication splitter for RecoverPoint, you must do the following:

1. Enable zero-padding on the Storage Pools, as described in [“Configuring the Installation Manager”](#).
 2. In the installation CSV file:
 - Add the **Splitter RPA IP** column, and enter the information, as described in [“Preparing the CSV topology file”](#).
 - In the **Optimize IOPS** column, set the value to `Yes`.
 3. Upload the splitter package to the IM server.
- Continue with [“Installing with the full Installation Manager”](#).

Preparing the IM on a Windows server

Prepare the Installation Manager (IM) on a Windows server.

Before you begin

Before deploying any ScaleIO components on a server that was part of another ScaleIO system, use the operating system tools to remove all of those ScaleIO components.

You can use a Windows IM to deploy to Windows and Linux servers.

Don't install the Gateway on a server on which RFcache (the xcache package) will be installed.

Procedure

1. Download and extract the files for your operating system.

You can download all files from the Complete Software download for your software version, available from Online Support (https://support.emc.com/products/33925_ScaleIO-Software).

The following table lists the files:

Required files	<ul style="list-style-type: none"> • Core files. You need these files for all operating systems of nodes in the system. <ul style="list-style-type: none"> ▪ LIA ▪ MDM ▪ SDC ▪ SDS • Gateway
Optional files	<ul style="list-style-type: none"> • GUI • XCACHE (for RFcache)

2. From the extracted download file, copy the ScaleIO Gateway MSI to the IM server:

```
EMC-ScaleIO-gateway-2.5-<build>.X-x64.msi
```

3. Run the file, and enter a new Gateway Admin password that will be used to access the IM.

The password must meet the following criteria:

- At least 8 characters long
 - Include at least 3 of the following groups: [a-z], [A-Z], [0-9], special chars (!@#\$...)
4. From the extracted download file, run the `Clean_XC_registry.bat` script on every Windows machine to be part of the system.
The script is included in the Complete Windows download.
 5. Prepare disks for storage:
 - a. Ensure that devices to be used are on-line, and initialized as MBR or GPT.
 - b. On each disk, use the Windows Disk Management to create a new, simple volume.
Assign a drive letter, and select **Do not format this volume**.
 6. Choose which IM mode to use:
 - To use the full IM, continue with [“Installing with the full Installation Manager”](#).
 - To use the IM wizard, skip to [“Installing with the Installation Manager wizard”](#).

Note

You can configure the Installation Manager, as described in [“Configuring the Installation Manager”](#)

Installing with the full Installation Manager

This section describes how to use the full Installation Manager to install ScaleIO components. To use the IM wizard, see the "Installing with the Installation Manager wizard" section of the *ScaleIO Deployment Guide*.

Note

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

To use the full Installation Manager, you must first prepare a CSV topology file.

Preparing the CSV topology file

Create a CSV topology file to deploy ScaleIO with the Installation Manager (IM).

Before you begin

Ensure that the settings you enter comply with ScaleIO product limits, as described in the "Product limits" table.

Note

This topic describes how to deploy ScaleIO on Linux and Windows nodes. Follow the instructions that match your environment and the ScaleIO support matrix for your version.

You can edit a CSV file with Excel or file-editing software. In this document, we will refer to and illustrate the CSV file as a spreadsheet.

The following CSV templates are provided as part of the software download, in the Gateway software packages:

- **Complete**

This spreadsheet template contains all available fields, both required and optional.

- **Minimal**

This spreadsheet template contains only the required fields. The optional fields (those that are in the complete spreadsheet, but not in the minimal one) will be assigned default values.

Procedure

1. Fill in your site-specific information in the appropriate places, overwriting the default information provided in the file.

You only need to use one spreadsheet for the installation, as follows:

- To manually enter all configuration details, use the Complete spreadsheet.
- To use default values for the non-mandatory fields, use the Minimal spreadsheet.
- To configure non-default values for columns that are not in the Minimal spreadsheet, either use the Complete spreadsheet, or copy the column heading from there into the minimal spreadsheet and enter your custom values into the minimal spreadsheet.

The following figure illustrates part of the Complete CSV file, to add Linux and Windows nodes:

Figure 12 CSV—complete

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R
1	Domain	Username	Password	Operating System	Is MDM/TB	MDM Mgmt IP	MDM IPs	MDM Name	perfProfile ForMDM	Virtual Ips	Virtual IP NICs	Is SDS	SDS Name	SDS All IPs	SDS-SDS C	SDS-SDS C	SDS-SDS C	SDS-SDS C
2		root	Password1	linux	Master	10.76.1.1	10.76.60.195	MDM1	Default	10.76.60.203	eth1	Yes		10.76.60.195			domain1	fs1
3		root	Password1	linux	Slave	10.76.1.2	10.76.60.17	MDM2	Default	10.76.60.203	eth1	Yes		10.76.60.17			domain1	fs1
4		root	Password1	linux	Slave	10.76.1.3	10.76.60.18	MDM3	Default	10.76.60.203	eth1	Yes		10.76.60.18			domain1	fs1
5		root	Password1	linux	TB		10.76.60.33	TB1	Default			Yes	SDS1	10.76.60.33	10.76.2.2,10.76.3.3		domain1	fs2
6	localhost	administrator	Password1	windows	TB		10.76.60.19	TB2	Default			No						
7		root	Password	linux								Yes-1	MultipleSDs_1	10.76.60.55			domain1	fs3
8		root	Password1	linux	Standby-Slave	10.76.1.4	10.76.60.20	MDM4	Default	10.76.60.203	eth1	No						
9	localhost	administrator	Password1	windows	Standby-TB		10.76.60.21	TB3	Default			No						
10		root	Password	linux								Yes-2	MultipleSDs_2	10.76.60.55			domain1	fs3
11		root	Password	linux								Yes-3	MultipleSDs_3	10.76.60.55			domain1	fs3
12		root	Password	linux								Yes-4	MultipleSDs_4	10.76.60.55			domain1	fs3

2. The following differences exist between Linux and Windows:

- a. For Linux nodes, the **Domain** column is not relevant. Leave the column blank or remove the column.
- b. Multiple SDS is supported on Linux nodes only.

The following table describes the fields in the spreadsheets. The required fields appear in both spreadsheets. Field names are case-sensitive; the order of the columns is not significant.

Table 9 CSV topology spreadsheets

Field	Description	Required
Domain	If using a domain user, the name of the domain (not relevant for Linux)	
Username	The name of the user. In Linux, this value is either <code>root</code> or a non-root sudo user, indicated by appending "(sudo)" to the end of the username. For more information on deploying with a non-root user, see Deployment of ScaleIO using a non-root user on page 280.	

Table 9 CSV topology spreadsheets (continued)

Field	Description	Required
	In Windows, this value is a user with administrator rights (default: <code>administrator</code>).	
Password	Password used to log in to the node. This should be the password of the user entered in the Username column. To authenticate with SSH instead of node passwords, see “Using SSH authentication on the ScaleIO Gateway” .	Yes
Operating System	The server’s OS: <code>linux</code> or <code>windows</code>	Yes
Is MDM/TB	The MDM role to deploy on this node: <code>Master</code> , <code>Slave</code> , <code>TB</code> , <code>Standby-Slave</code> , <code>Standby-TB</code> , or blank (if not an MDM). For more information, see "The MDM cluster."	
MDM Mgmt IP	The IP address for the management-only network	
MDM IPs	MDM IP addresses used to communicate with other ScaleIO components in the storage network. Maximum of eight addresses, comma-separated, no spaces.	
MDM Name	The name to be assigned to an MDM node. MDM names must be unique and may not contain spaces. Ensure that each server on which MDM is installed has a unique hostname.	
perfProfileForMDM	Optional performance profile to set for MDMs: <code>High</code> or <code>Default</code> (same as leaving this field empty).	
Virtual IPs	A virtual IP address for each possible manager MDM. This VIP can be used for communication between the MDM cluster and SDCs. Only one virtual IP address can be mapped to each NIC, with a maximum of four virtual IP addresses per system. Virtual IP addresses are not supported on nodes using Windows operating system.	
Virtual IP NICs	The NIC to which the virtual IP addresses are mapped.	
Is SDS	Deploy the SDS component on this node: <code>Yes</code> or <code>No</code> To configure multiple SDSs on a server, see “Configuring multiple SDSs on one server” . This is not supported on Windows nodes.	Yes
SDS Name	The name for the SDS node. The name cannot contain spaces.	
SDS All IPs	SDS IP addresses to be used for communication among all nodes. Maximum of eight addresses, comma-separated, no spaces.	
SDS-SDS Only IPs	SDS IP addresses to be used for communication among SDS nodes only. Maximum of eight addresses, comma-separated, no spaces.	
SDS-SDC Only IPs	SDS IP addresses to be used for communication among SDS and SDC nodes only. Maximum of eight addresses, comma-separated, no spaces. For SDC-only nodes, enter the IP in this column.	
Protection Domain	The Protection Domain to which to assign this SDS.	
Fault Set	The Fault Set to which to assign this SDS.	

Table 9 CSV topology spreadsheets (continued)

Field	Description	Required
	<p>Note</p> <p>When using Fault Sets, you must assign a minimum of three, and follow the guidelines described in “Fault Sets”.</p>	
SDS Storage Device List	<p>Storage devices to be added to an SDS. For more than one device, use a comma-separated list, with no spaces.</p> <p>Ensure that devices are prepared as described in "Configuring direct attached storage" in the Architecture section of the user documentation.</p> <p>Device name format:</p> <ul style="list-style-type: none"> Windows: <code>d,PhysicalDrive1,PhysicalDrive2,f,g</code> To add more devices than the amount of available drive letters, and for other options, see “Adding devices to SDS nodes on Windows servers”. Linux: <code>/dev/sdb,/dev/sdc</code> When specifying the SDS device path on a Linux node, use the path according to how it is listed in <code>cat/proc/partitions</code> (and not according to the output of <code>fdisk -l</code>) For example: <code>fdisk output: /dev/mapper/samah-lv1, /dev/sdb</code> <code>cat /proc/partitions output: dm-3, sdb</code> Use these values in the CSV: <code>/dev/dm-3, /dev/sdb</code> <p>To enable volume creation, you must add (at least) one device to (at least) 3 SDSs, where each SDS is in a separate fault unit, and each device has a minimum of 100 GB free storage capacity. You can do that via the CSV, or at a later stage. The maximum number of devices per SDS is listed in the “Product limits” table.</p> <p>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.</p> <p>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. After adding devices, you can define how much of the device capacity is available to ScaleIO by using the SCLI <code>modify_sds_device_capacity</code> command.</p> <p>When adding devices that were used in a previous ScaleIO system, follow the instruction in “Adding SDS devices that were used in a previous ScaleIO system”.</p> <p>Balance the total device capacity over all SDSs.</p>	
StoragePool List	<p>Sets Storage Pool names.</p> <p>Appropriates Storage Pools to the devices in the SDS Storage Device List, respectively.</p> <p>The amount of Storage Pools must equal the amount of devices. Comma-separated, no spaces.</p> <p>If no Storage Pool is set, one will be automatically created during installation, named <code>default</code>.</p>	
SDS Storage Device Names	Sets names for devices.	

Table 9 CSV topology spreadsheets (continued)

Field	Description	Required
	Appropriates device names to the devices in the SDS Storage Device List , respectively. The amount of names must equal the amount of devices. Comma-separated, no spaces.	
RFcache	<p>Install the xcache package on an SDS: <i>Yes</i> or <i>No</i> (default). If you are installing xcache on a Windows server, a server restart is necessary.</p> <hr/> <p>Note</p> <p>If you wish to configure RFcache devices as part of the deployment, set this column to <i>Yes</i>, RFcache Device List column to a list of comma-separated devices, and the RFcache Pool List to a corresponding comma-separated list of Storage Pools.</p>	
RFcache Device List	List of devices to provide RFcache acceleration on an SDS. Up to eight devices, comma-separated, with no spaces. If RFcache is <i>Yes</i> , and this field is left blank, you can add RFcache devices after installation.	
RFcache Pool List	<p>The name of the Storage Pools for which RFcache acceleration is to be provided. When using this field, the RFcache pool will be appropriated to the devices in the RFcache Device List, respectively. The amount of names must equal the amount of devices. Comma-separated, no spaces.</p> <p>If there are values for RFcache Device List, there must be values for this field, too. The specified Storage Pools must be defined in StoragePool List.</p>	
perfProfileForSDS	Optional performance profile to set for SDSs. <i>High</i> or <i>Default</i> (same as leaving this field empty).	
Is SDC	<i>Yes</i> or <i>No</i> . External SDC on RHEL 7.4 is supported on bare-metal servers only, not as guests on a hypervisor.	Yes
perfProfileForSDC	Optional performance profile to set for SDCs. <i>High</i> or <i>Default</i> (same as leaving this field empty).	
SDC Name	Sets name for the SDC.	

Adding SDS devices that were used in a previous ScaleIO system

Devices that were used in previous ScaleIO systems may not have been removed properly. When trying to add these device, the deployment will fail. This is most likely to occur in the following cases:

- When installing ScaleIO on SDS servers whose devices were previously used as SDS devices.
- When extending a system, with an SDS whose devices were previously used as SDS devices.
- When extending a system, by adding an additional SDS (multiple SDS) to a node that already has an SDS installed on it, regardless of the devices that are being used.

To prevent this issue, perform the following, before running the deployment:

Procedure

1. In the `gatewayUser.properties` file, set the value of `add.sds.with.force.on.next.run = true`.

For more information, see [“Configuring the Installation Manager”](#).

2. Restart the `scaleio-gateway` service.

After finishing and marking the deployment complete, the flag will revert to `false`, so you will need to set it again, as necessary.

Configuring multiple SDSs on one server

To configure multiple SDSs on one server, follow these guidelines:

- Multiple SDS is supported with the full installation only, not with the installation wizard.
- Multiple SDS is supported on Linux only.
- Use a separate line for each SDS. This will result in multiple lines for the same server.
- On any server, you cannot mix one SDS and multiple SDSs.
- In the **Is SDS** column, enter `yes-X` on each line, where *X* is the number of the SDS. For example: `Yes-1`, `Yes-2`

Note

On any SDS IP address, you should enter either one line with the `Yes` value, or a minimum of two lines with `Yes-X` options. You cannot mix the `Yes` value (which means single SDS) with any `Yes-X` value (which means multiple SDS) on the same IP address.

- When assigning multiple SDSs to a server, the SDS numbers must be consecutive, beginning with 1. For example, to assign three SDSs to one server:
 - Valid—`Yes-1`, `Yes-2`, and `Yes-3`
 - Invalid—`Yes-1`, `Yes-2`, and `Yes-4`
- When extending a node on which there are multiple SDS, you must first set `add.sds.with.force.on.next.run = true`, as described in [“Adding SDS devices that were used in a previous ScaleIO system”](#).

In the following example, the first four lines (under the field names) represent one server with four SDSs. The last line is a different server, with only one SDS.

D	E	F
Is SDS	SDS Name	SDS Device List
Yes-1	<code>sds_76-1</code>	<code>/dev/sdb,/dev/sdc</code>
Yes-2	<code>sds_76-2</code>	<code>/dev/sdd,/dev/sde</code>
Yes-3	<code>sds_76-3</code>	<code>/dev/sdf,/dev/sdg</code>
Yes-4	<code>sds_76-4</code>	<code>/dev/sdh,/dev/sdi</code>
Yes	<code>sds_244</code>	<code>/dev/sdb,/dev/sdc,/de</code>

Installing with the Installation Manager

This section describes how to install and configure ScaleIO components using the Installation Manager.

You need to do the following steps:

1. Log in to the IM server.
2. Upload installation packages.
3. Upload CSV file.
4. Initiate the installation.
5. Complete the installation.

Log in to the IM server

Procedure

1. Log in to: `https:// <IM_Server_IP>`
where `<IM_Server_IP>` is the IP address of the server where you installed the IM package.
2. Accept the certificate warning; alternatively, install your own certificate for the Tomcat server.
3. Enter the default user name, `admin`, and the password defined when the IM was prepared, then click **Login**.

The **Welcome** screen appears.

Upload installation packages

Procedure

1. Click **Packages**.

You may need to re-authenticate with your login credentials. The **Manage Installation Packages** window appears.

Note

To use this Gateway to install packages for Windows and Linux, you can upload all the files at once.

2. Browse to where the ScaleIO packages are located for the OS of the servers you are deploying.
3. Select all the relevant files for the operating systems you want to deploy on, then click **Open**.

File formats can be RPM, TAR, or MSI. Minimally, you must select these packages:

- a. MDM
- b. SDS (To enable multiple SDSs on one server, select all the SDS packages.)
- c. SDC
- d. LIA

The **Browse** button changes its appearance to reflect the selected files.

4. Click **Upload**.

The uploaded installation packages (RHEL, in this example) are listed in the file table.

5. Click **Proceed to Install**.

Note

To install the replication splitter for RecoverPoint, be sure to upload the splitter package. Replication support is version-specific; see the EMC Simple Support Matrix (ESSM) at <https://elabnavigator.emc.com/eln/elhome>.

The **Provide Installation Topology** screen appears.

Upload CSV file

If you have not created the CSV file yet, you can download a template by clicking **Minimal** or **Complete** on this screen.

Procedure

1. Click **Browse**, browse to where the installation CSV file is located, select the file, and click **Open**.
2. For a new installation, select **New installation**.

To extend an existing installation, click the down-arrow and select **Add to existing sys**.

3. Click **Upload Installation CSV**.

After successfully uploading the file, the **Installation Configuration** screen appears.

You can expand or contract the lists in the **Topology** section.

Initiate the full installation

Procedure

1. Enter a new MDM password.

The MDM password is a password that is configured for the MDM during the installation. It must meet the following criteria:

- Between 6 and 31 ASCII-printable characters
 - Includes at least 3 of the following groups: [a-z], [A-Z], [0-9], special chars (!@#\$...)
 - Contains no white spaces
-

Note

When using the Installation Manager to extend (rather than to install) a system, enter the MDM credentials that you entered during the initial installation.

2. Enter the LIA password.

The LIA password is a new password that will be used to authenticate communication between the IM and the LIA. It must meet the same criteria as the MDM password, as listed in the previous step.

Note

When using the Installation Manager to extend (rather than to install) a system, enter the LIA credentials that you entered during the initial installation.

3. Review and accept the end user license agreement.

Note

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 EMC Support at <https://support.emc.com>.

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

4. To use any of the following advanced options, select **Set advanced options**:

You can use the Installation Manager to configure Syslog event reporting. You can also configure these features after installation, via the CLI.

Option type	Option	Description
Skip options	Skip upload	Don't upload packages. You can use this if the packages were already uploaded.
	Skip installation	Don't install packages. You can use this if the packages were already installed.
	Skip configuration	Don't configure ScaleIO components. You can use this if you only want to upload and install packages.
	Skip Linux devices validation	Don't validate Linux device names. For more information, see the "Installing without validating Linux devices" section of the <i>ScaleIO Deployment Guide</i> .
Configuration options	Enable zero-padding	<p>Enable zero-padding of new Storage Pools.</p> <p>Zero padding is required for using the background scanner in data comparison mode and for use with RecoverPoint splitter replication.</p> <hr/> <p>Note</p> <p>Replication support is version-specific; see the ESSM for full details.</p>
	Enable alert service	<p>Enable the alert service required for SNMP and ESRS reporting. This also creates and configures the lockbox.</p> <p>If you select this option, the Traps Receiver IP/</p>

Option type	Option	Description
		Hostname field is displayed. Enter up to two (comma-separated) IP addresses or hostnames of the SNMP trap receiver servers.
Security options	Disable secure communication with MDM	<p>Disable the need for secure communication mode between management clients and the MDM.</p> <hr/> <p>Note</p> <p>Disabling secure communication has security implications, described in “Using SCLI in non-secure mode”.</p> <hr/>
	Disable secure communication with LIA	<p>Disable the need for secure communication mode between management clients and the LIA.</p> <hr/> <p>Note</p> <p>Disabling secure communication has security implications, described in “Using SCLI in non-secure mode”.</p> <hr/>
	Disable authentication in internal components	Disable the need for authentication between SDSs and MDMs.
	Use trusted IP addresses	<p>Limit the ScaleIO Gateways that can communicate with this LIA.</p> <p>If you select this option, enter the IP addresses to allow, including the IP address of this Gateway. If the Gateway uses multiple IP address (for example, one for management and another for data), enter all the addresses.</p>

5. To configure Syslog reporting, select **Configure MDM to send messages to syslog server**, and enter the following parameters:

- a. **Syslog Server**

The host name or IP address of the syslog server to where the messages are to be sent. Enter up to two servers, comma-separated.

b. Port

The port of the syslog server (default 1468)

c. Syslog Facility

The facility level (default: Local0)

- Review the displayed information, then click **Start Installation**.

Note

If you are installing RFcache (the `xcache` package) on a Windows server, a server restart is necessary. To continue with the installation, click **Allow restart**.

- Click the **Monitor** tab.

Complete the installation

The installation performs the following phases: query, upload, install, and configure. The monitor **Install - query** screen appears:

The screenshot shows the 'Monitor' tab in the ScaleIO interface. The breadcrumb navigation is: Home > Packages > Install > **Monitor** > Maintain > Logout. The main heading is 'Operation Progress' with a sub-heading 'Server: install - query phase'. There are three buttons: 'Start upload phase', 'Abort', and 'Rerun phase'. Below this, a summary bar shows: 'Progress of this phase: running: 0, completed: 5, pending: 0, aborted: 0, failed: 0 (10:24:57 AM)'. A table lists the progress of five 'Query' phases, all completed. The table has columns: Phase, IP, Command, Status, Start time, and Details. At the bottom, there are checkboxes for 'Auto Refresh' (checked), 'Hide Completed Successfully' (unchecked), and 'Show All Phases' (unchecked). A pagination bar shows 'Page 1 of 1' and 'Viewing 1 - 5 of 5'.

Phase	IP	Command	Status	Start time	Details
Query	Installation Manager	validate and orchestrate new command...	✓ completed	12 minutes ago	Details
Query	192.168.1.229	validate node	✓ completed	12 minutes ago	Details
Query	10.103.110.246	validate node	✓ completed	12 minutes ago	Details
Query	10.103.110.239	validate node	✓ completed	12 minutes ago	Details
Query	10.103.110.230	validate node	✓ completed	12 minutes ago	Details

In the query stage, the IM validates that there are no previously installed ScaleIO components on any of the requested installation nodes.

When extending an existing system, the query phase does expect to find currently-installed nodes.

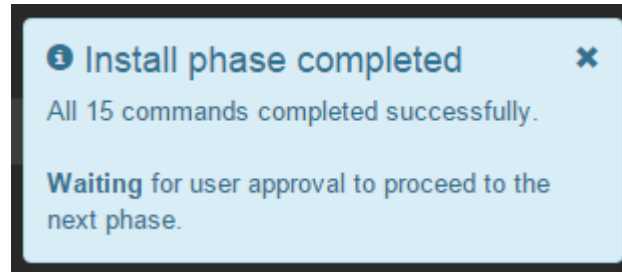
Note

When using the IM to install the replication splitter for RecoverPoint, this step is not performed on the SDC nodes. Replication support is version-specific; see the ESSM for full details.

You can change the display by selecting the following:

- **Auto Refresh** (default)
- **Hide Completed Successfully**
- **Show All Phases**

At the end of each phase, a status message, similar to the following is displayed. You must approve moving to the next phase.



If an error message is displayed during the process, you can retry the failed operations (with the **Retry failed** button), continue to the next phase, or abort the installation.

Note

Once you proceed to the next stage, you will not be able to use the IM to retry failed tasks.

Procedure

1. When the query phase is complete, click **Start upload phase**.
The **Install - upload** screen appears, displaying which ScaleIO packages are being uploaded to each server, and the status of each command.
2. When the previous phase is complete, click **Start install phase**.
The **Install - install** screen appears, displaying the status of the installation commands.
3. When the previous phase is complete, click **Start configure phase**.
The **Install - configure** screen appears, displaying the status of the configuration commands.
4. When all processes are finished, click **Mark operation completed**.
The ScaleIO system installation is complete!

Note

Marking the operation completed signals to the IM that it can now be used for other installations, upgrades, and so on.

A post-operation notice is displayed. The steps in this notice are described in the post-deployment checklist in the *EMC ScaleIO Deployment Guide*, as well as in the *EMC ScaleIO User Guide*.

After you finish

It is highly recommended to run the ScaleIO system analysis to analyze the ScaleIO system immediately after deployment, before provisioning volumes, and before using the system in production. For more information, see "System analysis overview."

Best practice suggestion

For optimal performance in environments with more than 60,000 IOPS, see the *ScaleIO Performance Fine-Tuning Technical Notes*.

ScaleIO support on Linux servers running Veritas Cluster Server (VCS) software

When a machine running VCS software on top of ScaleIO volumes is rebooted, the VCS virtual devices may enter failed state. This is due to a change introduced in

ScaleIO v2.0, that can cause the SDC driver to temporarily hang, and the ScaleIO volumes to appear only after a 45-second delay during machine boot. When the VCS devices enter the failed state, manual intervention is required to refresh them, using VCS command line utilities.

To bypass this issue, perform the following steps on every server running the SDC, after installing the SDC:

Procedure

1. Open `/etc/init.d/scini` for editing.
2. In the line that calls `drv_cfg --boot_wait`, add this text: `--seconds 2`

The corrected line should read:

```
/opt/emc/scaleio/sdc/bin/drv_cfg --boot_wait --seconds 2
```

3. Save the edited file.

Results

This change shortens the SDC driver's wait time on the MDM during machine reboot and (in most of the cases) will avoid the hang. In the very rare case where the SDC didn't receive its first MDM update in 2 seconds, the ScaleIO volumes will appear only after 45 seconds, causing the VCS devices to fail as described above. If this happens, use VCS command line utilities to rediscover the SDC volumes.

Where to go from here

Now that your ScaleIO system is up and running, the next task is to create and map volumes. The SDCs expose these volumes as local storage devices to the applications servers. For more information, see the "SAN virtualization layer" section in the ScaleIO documentation.

You can create and map volumes using any of the ScaleIO CLI management utilities, as described in the *EMC ScaleIO User Guide*.

Installing with the Installation Manager wizard

This section describes how to use the installation wizard, the quickest way to get a ScaleIO system up and running.

Note

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 EMC Support at <https://support.emc.com>.

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

The wizard is most suitable for environments when all servers are converged and in one Protection Domain.

Note

To use the wizard, you must prepare the Installation Manager, as described in [“Preparing the Installation Manager and the Gateway”](#).

The wizard prepares all ScaleIO management and data communication on a single network. The IM must have access to that network to succeed. You can add additional data network IP addresses for SDS-SDS and SDS-SDC communication via the CLI or the GUI.

The wizard guides you through the following tasks:

1. Setting up the installation
2. Running the installation

Setting up the installation

This section describes how to set up the wizard installation:

Procedure

1. Log in to the Installation Manager (IM) server:
 - a. Point your browser to this URL: `https:// <IM_Server_IP>`
 where `<IM_Server_IP>` is the IP address of the server where you installed the IM package
 - b. Accept the certificate warning; alternatively, install your own certificate for the Tomcat server.
 - c. Enter the default user name, `admin`, and the password defined when the IM was prepared, then click **Login**.
2. Upload installation packages:
 - a. From the **Welcome** screen, click **Packages**.
 You may need to re-authenticate.
 - b. From the **Manage Installation Packages** screen, browse to the extracted ScaleIO component packages, select the files, and click **Open**.
 File formats can be RPM, TAR, or MSI.
 Minimally, you must select these packages:
 - MDM
 - SDS (The SDS-X packages are not needed.)
 - SDC
 - xcache
 - LIA
 - c. Click **Upload**.

The uploaded installation packages appear in the file table.

ScaleIO Installer EMC² ScaleIO

Home **Packages** Install Monitor Maintain Logout

Manage Installation Packages

Successfully uploaded

Select the RPM / TAR / MSI packages to install on your nodes. They will be distributed to the nodes in the next step.
A new installation requires the following packages: MDM, SDS, SDC, LIA and XCACHE (optional).

Browse ... Upload Delete

Packages uploaded to Installation Manager

Type	OS	Linux flavor	Version	Latest	Size	File name
SDS3	Linux	RHEL6	2.0-10000.2012	✓	4.66 MB	EMC-ScaleIO-sds3-2.0-10000.2012.el6.x86_64.rpm
SDS1	Linux	RHEL6	2.0-10000.2012	✓	4.65 MB	EMC-ScaleIO-sds1-2.0-10000.2012.el6.x86_64.rpm
SDC	Linux	RHEL6	2.0-10000.2012	✓	6.09 MB	EMC-ScaleIO-sdc-2.0-10000.2012.el6.x86_64.rpm
SDS	Linux	RHEL6	2.0-10000.2012	✓	4.67 MB	EMC-ScaleIO-sds-2.0-10000.2012.el6.x86_64.rpm
SDS2	Linux	RHEL6	2.0-10000.2012	✓	4.65 MB	EMC-ScaleIO-sds2-2.0-10000.2012.el6.x86_64.rpm
MDM	Linux	RHEL6	2.0-10000.2012	✓	8.81 MB	EMC-ScaleIO-mdm-2.0-10000.2012.el6.x86_64.rpm
SDS4	Linux	RHEL6	2.0-10000.2012	✓	4.65 MB	EMC-ScaleIO-sds4-2.0-10000.2012.el6.x86_64.rpm
LIA	Linux	RHEL6	2.0-10000.2012	✓	3.98 MB	EMC-ScaleIO-lia-2.0-10000.2012.el6.x86_64.rpm
XCACHE	Linux	RHEL6	2.0-10000.2012	✓	590.59 KB	EMC-ScaleIO-xcache-2.0-10000.2012.el6.x86_64.rpm

Page 1 of 1 10 Viewing 1 - 9 of 9

Show only latest packages Proceed to Install >

d. Click **Proceed to Install**.

3. In the **Provide Installation Topology** screen, click the **Installation Wizard** drop-down arrow.

4. Select the format of the MDM cluster:

- 3-node cluster (default)

Configures a Master MDM, a Slave MDM, and a Tie Breaker MDM.

- 5-node cluster

Configures a Master MDM, two Slave MDMs, and two Tie Breaker MDMs.

The **Installation Configuration** window appears:

Home Packages **Install** Monitor Maintain Logout

Installation Configuration

Review the configuration, enter passwords and other options, then start the installation.

Ensure that OpenSSL is installed on all nodes in the system. For Windows nodes use the following links:

- [C++ redistributable](#)
- [OpenSSL for 64bit](#)

Credentials Configuration

MDM Password: ?

LIA Password: ?

License: ☒ I accept the terms of the End User License Agreement (mandatory). [View EULA](#)

Topology

Click on a table cell to edit (except for the role which is fixed).
Press <ENTER> to apply changes.

Role	IP	OS	MDM Name	Password
MDM Master / SDS / SDC	10.0.0.1	Linux	Manager1	
MDM Slave / SDS / SDC	10.0.0.2	Linux	Manager2	
MDM Slave / SDS / SDC	10.0.0.3	Linux	Manager3	
Tie Breaker / SDS / SDC	10.0.0.4	Linux	Tie-Breaker1	
Tie Breaker / SDS / SDC	10.0.0.5	Linux	Tie-Breaker2	

Viewing 1 - 5 of 5

+ Clone host ✕ Delete host ☒ Show passwords while editing

Start Installation > - OR - ✕ Cancel

Initiate the wizard installation

Start the deployment of ScaleIO.

Procedure

1. Enter a new MDM Password and LIA Password.

The MDM and LIA passwords must meet the following criteria:

- Between 6 and 31 ASCII-printable characters
- Include at least 3 of the following groups: [a-z], [A-Z], [0-9], special chars (!@#\$...)
- No white spaces

The **Topology** section is populated with nodes, according to the MDM cluster chosen in the previous step. Each node is also defined as an SDS and an SDC.

2. Review and accept the End User License Agreement (EULA).

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 EMC Support at <https://support.emc.com>.

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

3. In the **Topology** section, enter server information:
 - a. For each node, change the IP address, select the host operating system, and enter the host password.

- For Linux, type the password of the root user.
- For Windows, type the password of an administrator user.
- To see the passwords, select **Show passwords while editing**.

b. To add more hosts, click **Clone host**.

The assigned IP is derived from the IP of the previous host. Each cloned host is enabled to act as an SDS and an SDC.

4. Click **Start Installation**.

Note

A post-installation notice appears, that outlines the steps that are required after the wizard completes. The content of this notice will be repeated at the end of the process, or you can view it by clicking **Post installation instructions** from the **Monitor** screen.

Running the installation

The installation performs the following phases: query, upload, install, and configure. At the end of each phase, you need to approve moving to the next phase.

Procedure

1. Click the **Monitor** tab.

The **Install - query** screen appears.

Home Packages Install **Monitor** Maintain Logout

Operation Progress
Server: install - query phase

Start upload phase Abort Rerun phase

Progress of this phase: running: 0, completed: 5, pending: 0, aborted: 0, failed: 0 (10:24:57 AM)

Phase	IP	Command	Status	Start time	Details
Query	Installation Manager	validate and orchestrate new command...	✓ completed	12 minutes ago	Details
Query	192.168.1.229	validate node	✓ completed	12 minutes ago	Details
Query	10.103.110.246	validate node	✓ completed	12 minutes ago	Details
Query	10.103.110.239	validate node	✓ completed	12 minutes ago	Details
Query	10.103.110.230	validate node	✓ completed	12 minutes ago	Details

Page 1 of 1 10 Viewing 1 - 5 of 5

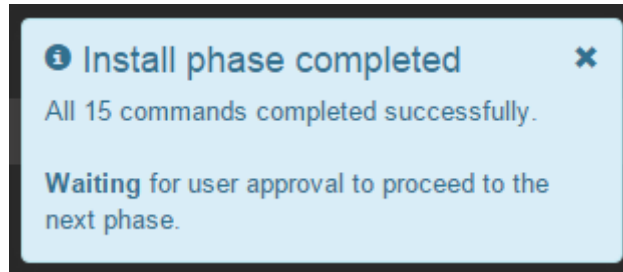
☒ Auto Refresh ☐ Hide Completed Successfully ☐ Show All Phases

In the query stage, the IM validates that there are no previously installed ScaleIO components on any of the requested installation nodes.

You can change the display by selecting the following:

- **Auto Refresh** (default)
- **Hide Completed Successfully**
- **Show All Phases**

At the end of each phase, a status message, similar to the following is displayed. You must approve moving to the next phase.



If an error message is displayed during the process, you can retry the failed operations (with the **Retry failed** button), continue to the next phase, or abort the installation.

Note

Once you proceed to the next stage, you will not be able to use the IM to retry failed tasks.

2. When the query phase is complete, click **Start upload phase**.

The **Install - upload** screen appears, displaying which ScaleIO packages are being uploaded to each server, and the status of each command.

3. When the previous phase is complete, click **Start install phase**.

The **Install - install** screen appears, displaying the status of the installation commands.

4. When all Install command rows are in Completed status, click **Start configure phase**.

The **Install - configure** screen displays configuration progress.

5. When all processes are finished, click **Mark operation completed**.

The ScaleIO system installation is complete! The wizard installation creates one Protection Domain and one Storage Pool, both named "default". These are used in the following section.

The post-installation notes appear, directing you to the steps necessary to start using your storage. These steps are described in the next section. It is highly recommended to run the ScaleIO system analysis to analyze the ScaleIO system immediately after deployment, before provisioning volumes, and before using the system in production. For more information, see [System analysis overview](#) on page 240.

After you finish

If you are using Linux servers that are running Veritas Cluster Server (VCS) software, see [ScaleIO support on Linux servers running Veritas Cluster Server \(VCS\) software](#) on page 83 for post-installation instructions.

Enabling the storage

After the ScaleIO system is installed, follow these steps to enable using the storage. You must issue these commands from the Master MDM, either directly, or via SSH/RDP.

Note

You can also use the GUI to add SDS devices.

Procedure

1. Add SDS devices, according to the following guidelines:
 - You must add at least one device per SDS, to at least 3 SDSs, with a minimum of 100 GB free storage capacity per device.
 - Balance the total device capacity among all SDSs.

- a. Log in, by typing the following command:

```
scli --login --username
<mdm_username> --password <mdm_password>
```

Note

If bash completion isn't enabled, run `./etc/bash_completion/scli`.

- b. Add devices, by typing the following command:

```
scli --add_sds_device --sds_ip <IP> --storage_pool_name
default --device_path <PATH>
```

Linux example:

```
scli --add_sds_device --sds_ip 192.168.212.10 --
storage_pool_name default --device_path /dev/sdX
```

Note

Mapped volumes appear to the SDC as `/dev/sciniX`, where X is a letter, starting from "a". For more information, see ["Mounting ScaleIO"](#).

Windows example:

```
scli --add_sds_device --sds_ip 192.168.212.10 --
storage_pool_name default --device_path d
```

Note

To add devices on a Windows-based SDS, the device path is either a drive letter or `PhysicalDriveX`. For more information, see ["Preparing devices on Windows servers."](#)

2. Add and map a volume:

a. Add a volume, by typing the following command:

```
scli --add_volume --storage_pool_name default --size_gb
<SIZE> --volume_name <NAME> --protection_domain <NAME>
```

Example:

```
scli --add_volume --storage_pool_name default --size_gb 16
--volume_name vol01 --protection_domain pd01
```

b. Map a volume to an SDC, by typing the following command:

```
scli --map_volume_to_sdc --volume_name <NAME> --sdc_ip <IP>
```

Example:

```
scli --map_volume_to_sdc --volume_name vol01 --sdc_ip
192.168.212.19
```

Your storage is ready to use.

To view the installed nodes and storage, use the GUI or the SCLI `--query_all` command.

Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers

Steps that must be taken when deploying SDC or RFCache on CoreOS, hLinux, Oracle Linux (OL), or Ubuntu servers.

Before you begin

Ensure that the ScaleIO SDC and RFCache (RF) drivers are compatible with the Ubuntu/OL/CoreOS kernels and that the SDC driver is compatible with the hLinux kernel.

Note

All references to Ubuntu in this section apply equally to supported versions of hLinux, OL, and CoreOS. For hLinux, this is relevant for SDC installation, not RFCache.

To ensure that a client node always has an SDC/RF kernel driver that matches the client kernel, the SDC/RF startup script uses `driver_sync.sh`, a synchronization utility that fetches SDC/RF kernel drivers from a configured repository. This repository can be either EMC's repository or a local repository in your private network.

For Ubuntu nodes with ScaleIO SDC/RF drivers:

- The `driver_sync.sh` script is located, by default, at `/bin/emc/scaleio/`.
- In the repository, a driver package tarball is located at `<base_dir>/Ubuntu/<scaleio_version>/<kernel_version>/scini.tar`.

For CoreOS nodes with ScaleIO SDC/RF drivers:

- The `driver_sync.sh` script is located, by default, at the location of the script in the OEM folder.
- In the repository, a driver package tarball is located at `<base_dir>/CoreOS/<scaleio_version>/<kernel_version>/scini.tar`.

It may be necessary to manually restart the driver/s and services on SDCs.

In a normal work flow, an SDC/RF driver will be loaded on boot. If there is a mismatch between the version of the running kernel and the kernel against which the driver was compiled, the `driver_sync.sh` script will be run automatically to fetch the appropriate driver. If the script fails for any reason (incorrect configuration etc.) after fixing the issue with the `driver_sync.sh` script, the driver services must be restarted manually, as described below. Root permission is required to run the commands.

To manually restart driver/s and services, run these commands:

```
service scini restart
```

```
service xcache restart
```

Ensuring the kernel version is correct

Before you begin, perform the following:

Procedure

1. Ensure that your kernel version matches a version in the repository, by performing the following:

Note

If all Ubuntu servers have the same kernel, you can run this command on a single server.

- a. Run this command on every Ubuntu server:

```
uname -r
```

Output similar to the following is displayed:

```
3.16.0-62-generic
```

- b. Copy the following text into a browser, and compare the output to the Ubuntu kernel version in the EMC repository:

```
ftp://QNzgdXix:Aw3wFAwAq3@ftp.emc.com
```

2. To use a mirror repository, you must ensure that the SSH public and SSH private keys are located in all system nodes in the same path.

The private key should have directory-level permission only (`chmod 700 <private_key_path>`). The public key can have all permissions. This is not necessary when using the EMC repository.

3. The GPG key must be located in all system nodes in the same path. This key is required for using a mirror or EMC repository.

After you finish

The configuration of `driver_sync.sh` is specified in `driver_sync.conf` file, which can be created in the following methods:

- [“Creating the configuration file via the ScaleIO Gateway”](#)
- [“Creating the configuration file manually”](#)

Creating the configuration file via the ScaleIO Gateway

This section describes how to create the `driver_sync.conf` file during the Gateway deployment process. Before deploying, you set parameters that will be used by the Gateway, during deployment, to create the file for your environment.

To edit the properties, perform the following:

Procedure

1. Use a text editor to open the `gatewayUser.properties` file, located in the following directory on the Gateway server:
`/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes`
2. Edit the file by adding the parameters described below.
3. Save and close the file.
4. Import the GPG key by running the following command, on every Ubuntu node that will contain SDC or RFcache:

```
apt-key add <new_public_key>
```

5. Restart the `scaleio-gateway` service by running the following command:

```
service scaleio-gateway restart
```

After the changes are made, use the Installation Manager to deploy the system, using the installation files for Ubuntu.

Parameters to add

- `sdc.kernel.repo.SCINI_REPO_ADDRESS`

The value of this variable is encoded as `<protocol>://[<address>]/<path>`

Where `<protocol>:= ftp | sftp | file`

The following example uses the address of the EMC repository:

```
sdc.kernel.repo.SCINI_REPO_ADDRESS=ftp://
QNzgdXix:Aw3wFAwAq3@ftp.emc.com
```

- `sdc.kernel.repo.SCINI_REPO_USER`

Contains the user name that is used as login name. Needed when the protocol is FTP or SFTP.

Example:

```
sdc.kernel.repo.SCINI_REPO_USER=admin
```

- `sdc.kernel.repo.SCINI_REPO_PASSWORD`

Represents the password. Needed only for FTP protocol.

Example:

```
sdc.kernel.repo.SCINI_REPO_PASSWORD=password
```

- `sdc.kernel.repo.SCINI_REPO_USER_KEY=<path_to_private_key>`

Contains a path to the private RSA key used for SSH connections (user identity file). Needed only for SFTP.

Example:

```
sdc.kernel.repo.SCINI_REPO_USER_KEY=<path_to_private_key>
```

- `sdc.kernel.repo.SCINI_REPO_HOST_KEY=<path_to_public_key>`

Contains a path to the repository server's public host key, used for SSH connections. Needed only for SFTP.

Example:

```
sdc.kernel.repo.SCINI_REPO_HOST_KEY=<path_to_public_key>
```

- `sdc.kernel.repo.SCINI_REPO_MODULE_SIGCHECK`

Set to 0 or 1, to determine whether the fetched kernel modules must be verified.

Example:

```
sdc.kernel.repo.SCINI_REPO_MODULE_SIGCHECK=1
```

- `sdc.kernel.repo.SCINI_REPO EMC_KEY`

Contains a path to the EMC GPG public key file that is needed for signature verification. You can retrieve the GPG key (RPM-GPG-KEY-ScaleIO_<version>.XXX.X.0) from the folder in the root of the ISO: ScaleIO_<version>.X.X_GPG-RPM-KEY_Download.zip

This is the same file used by LIA to verify RPMs. It needs to be set only if SCINI_REPO_MODULE_SIGCHECK was set to 1.

Example:

```
sdc.kernel.repo.SCINI_REPO EMC_KEY=<path_to_key>/RPM-GPG-KEY-ScaleIO_3.0.168.0
```

You can now deploy the system as described in "Installing with the full Installation Manager" in the *ScaleIO Deployment Guide*.

During the deployment, a `driver_sync.conf` file is generated for the SDC and for RFcache in their respective `bin` folders:

- SDC: `/bin/emc/scaleio/scini_sync`

- RFcache: /bin/emc/scaleio/xcache_sync

Creating the configuration file manually

This topic describes how to create the `driver_sync.conf` file manually. Actually, if you already deployed the system with the Gateway, the file already exists in the `bin` folder listed above, with dummy values, and you will need to edit that file, and then copy it to the `bin` folder on every SDC/RFcache (RF) server in the system.

The configuration of the `driver_sync` script is read from `driver_sync.conf` (an example appears below), which contains the following parameters:

Table 10 `driver_sync.conf` parameters

Parameter	Description
<code>repo_address</code>	Encodes the chosen method and address.
<code>repo_user</code>	Identifies the user name used when logging in to the drivers server. Required for SFTP and FTP methods.
<code>repo_password</code>	The FTP server login password (for user <code>repo_user</code>). Required for FTP method
<code>user_private_rsa_key</code>	Contains the path of the private RSA key file used to SFTP connection. This key file (provided by ScaleIO on installation) enables SFTP connection without a password.
<code>local_dir</code>	The local directory in which the downloaded SDC/RF drivers are stored. For SDC/RF clients, this value is set up during installation and doesn't normally require a change.
<code>repo_public_rsa_key</code>	The path of the public RSA key of the repository server machine (also known as the host public key).
<code>module_sigcheck</code>	Determines whether to verify (1) or not to verify (0) the downloaded kernel drivers.
<code>emc_public_gpg_key</code>	A key that is used to verify signatures of downloaded drivers. Enter the path to the public GPG signing key of ScaleIO (just like the one that comes with ScaleIO LIA).
The following parameters are used when <code>driver_sync.sh</code> is used to synchronize a local repository against a remote one.	
<code>sync_pattern</code>	A regular expression that can be used to fetch several drivers at once. This parameter is not normally needed, and becomes relevant only when <code>driver_sync.sh</code> is used by your own private repository to fetch from a remote (e.g. EMC) repository.
<code>is_local_repo</code>	Set to 0 (default) or omit, for SDC/RF clients. Set to 1 when <code>driver_sync.sh</code> is used to synchronize local repository against a remote repository. Setting to 1 will cause the fetched tar files not to be extracted.

Example:

```
#####
#driver_sync Configuration file
#Everything after a '#' until the end of the line is ignored
```

```
#####
#Repository address, prefixed by protocol
repo_address = sftp://localhost/path/to/repo/dir
#repo_address = ftp://localhost/path/to/repo/dir
#repo_address = file://local/path/to/repo/dir
# Repository user (valid for ftp/sftp protocol)
repo_user = scini
# Repository password (valid for ftp protocol)
repo_password = scini
# Local directory for modules
local_dir = /bin/emc/scaleio/scini_cache/
# User's RSA private key file (sftp protocol)
user_private_rsa_key = /bin/emc/scaleio/scini_key
# Repository host public key (sftp protocol)
repo_public_rsa_key = /bin/emc/scaleio/scini_repo_key.pub
# Should the fetched modules' signatures be checked [0, 1]
module_sigcheck = 1
# EMC public signature key (needed when module_sigcheck is 1)
emc_public_gpg_key = /bin/emc/scaleio/emc_key.pub
# Sync pattern (regular expression) for massive retrieve
sync_pattern = .*
```

Update the ScaleIO signature key

If the ScaleIO package-signing key expires, you must update the configuration for each SDC or RFcache node.

When a ScaleIO package-signing key expires, a new public key is issued. You must update `driver_sync.conf` to use the new public key.

Procedure

1. Download the new public key from the Customer Support site (<https://support.emc.com>).
2. On each SDC or RFcache node, edit `driver_sync.conf` so that the value of `emc_public_gpg_key` is the path to the new key file.

The expired public key will still be in the gpg keyring, and older packages that were signed with the old key can still be verified.

Location of `driver_sync.conf`

- SDC

Host	Path
Ubuntu/OL	/bin/emc/scaleio/scini_sync/driver_sync.conf
CoreOS	/usr/share/oem/bin/emc/scaleio/scini_sync/ driver_cache.conf

- RFcache

Host	Path
Ubuntu/OL	/bin/emc/scaleio/xcache_sync/driver_sync.conf
CoreOS	/usr/share/oem/bin/emc/scaleio/xcache_sync/ driver_sync.conf

Running the ScaleIO Gateway on CoreOS container

Run the ScaleIO Gateway on the docker container

This step is necessary when you do not have another Gateway available to deploy ScaleIO.

When building the Dockerfile, keep the following in mind:

- Set the base image to CentOS 6.x.
- Install Java 1.8.
- Expose the ports necessary for the Gateway (port 80 and port 443).

Procedure

1. Use a Dockerfile to build a container with the following RPM:

```
EMC-ScaleIO-gateway-2.5-<build>.X.x86_64.rpm
```

2. Run the container.

Creating a mirror repository

This section describes how to set up a local repository, necessary when connecting all SDC/RFcache (RF) nodes to the ScaleIO repository is not possible, perhaps due to security considerations or other limitations.

A common workflow would be to configure your mirror repository to synchronize with the ScaleIO repository, while your SDC/RF nodes will fetch drivers (synchronize) from the mirror repository.

Driver repository directory hierarchy

The directory hierarchy in the repository consists of the following levels:

```
<Distro>/<ScaleIOVersion>/<Kernel Version>/
```

where:

- *<Distro>* is either CoreOS or Ubuntu or OEL (for Oracle Linux)
- *<ScaleIOVersion>* is encoded as x.x.x.x
- *<Kernel Version>* is the output of `uname -r`.

The files themselves are tar files named `scini.tar` and `xcache.tar`.

For example:

```
<Repo_base>/CoreOS/2.0.30.0/4.1.7-coreos/scini.tar is a tar file of the
SDC/RFcache driver for ScaleIO version 2.0.30.0 for CoreOS with kernel version
4.1.7-coreos.
```

Synchronizing the repository

In normal workflow, you do not have to (and should not) modify the directory structure manually. Instead, you should use `driver_sync.sh`, which provides a proper `driver_sync.conf`, as described in the previous section. With the configuration file properly written, synchronizing is done by simply calling `driver_sync.sh sync`.

In the `driver_sync.conf` file, you must set the `is_local_repo` parameter to 1. This will instruct `driver_sync.sh` not to untar the downloaded files.

It is also recommended that you change the `sync_pattern` parameter, a regular expression, to something more specific than `".*"`. Otherwise the sync operation will mirror the entire EMC driver repository every time.

Setting up a new SFTP-based repository

To save time in setting up a new SFTP-based driver repository, use the `driver_repo_wizard.sh` script. This script can be found in the following locations: under the `scini_sync` or `xcache_sync` directory in `/bin/emc/scaleio` (Ubuntu) or `/usr/share/oem/bin/emc/scaleio` (CoreOS).

- Ubuntu or Oracle Linux:
 - `/bin/emc/scaleio/scini_sync`
 - `/bin/emc/scaleio/xcache_sync/`
- CoreOS:
 - `/usr/share/oem/bin/emc/scaleio/scini_sync`
 - `/usr/share/oem/bin/emc/scaleio/xcache_sync`

The script does the following:

1. Creates a new user with a name you provide. The user will be created shell-less.
2. Creates a pair of RSA keys that will be used for passwordless SSH connection. You must provide the created private key to all your client nodes, while the public key will be added to the `authorized_keys` file of the created user.
3. Updates the SSH service configuration file (`/etc/ssh/ssh_config`) for the new user, instructing to `chroot` the user to its home directory, which should also be the repository base directory.

Deploying on OpenStack

This section describes how to deploy ScaleIO in the OpenStack environment.

The following versions are supported:

- Liberty
- Mitaka
- Newton
- Red Hat OpenStack Platform
- Mirantis OpenStack
- Canonical Ubuntu OpenStack.

ScaleIO with Mirantis OpenStack

EMC has developed open-sourced Fuel Plug-ins to simplify ScaleIO deployment as well as integrate your Mirantis OpenStack deployment with existing ScaleIO storage systems. Searching for ScaleIO on the Mirantis Fuel Plugin site provides links to software and documentation.

You can install ScaleIO manually, as well, or in an automated fashion with other automation frameworks. Visit <http://emccode.com> to see what open-sourced code exists to assist in automating ScaleIO.

ScaleIO with Canonical (Ubuntu) OpenStack

EMC has developed open-sourced Juju Charms to simplify deployments of ScaleIO as well as integrate your Canonical (Ubuntu) OpenStack deployment with existing ScaleIO storage systems. Searching for ScaleIO on the Canonical Juju Charm site provides links to software and documentation.

You can install ScaleIO manually, as well, or in an automated fashion with other automation frameworks. Visit <http://emccode.com> to see what open-sourced code exists to assist in automating ScaleIO.

ScaleIO with OpenStack Liberty

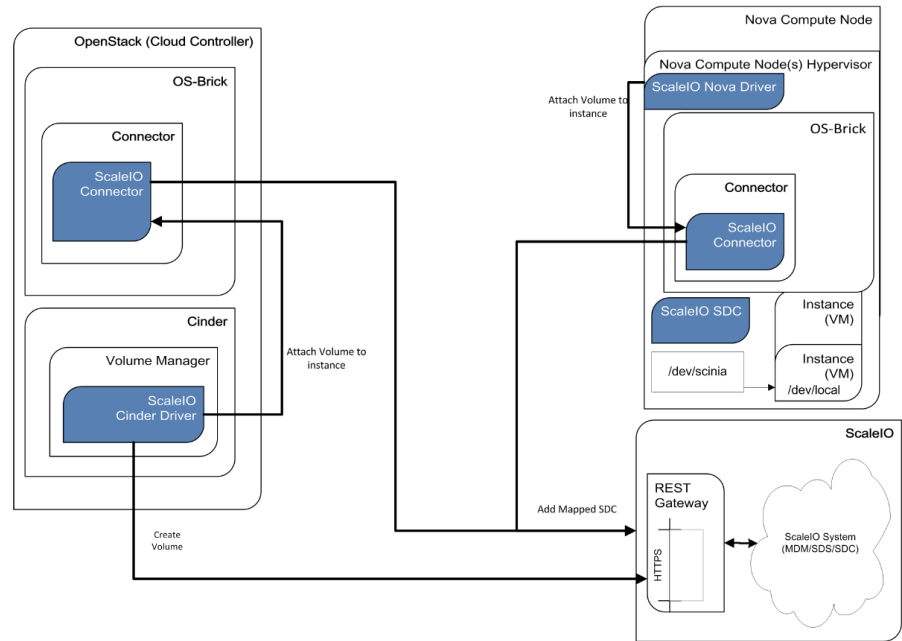
The ScaleIO elastic storage solution for OpenStack Liberty does not contain an OpenStack installation package, because the ScaleIO drivers are already included in the Liberty version. For information on configuring the driver, see the OpenStack documentation for the version.

The figure below illustrates the implementation of ScaleIO with OpenStack Liberty.

- Communication between the Cinder driver and the ScaleIO MDM is implemented using the REST API. This happens usually in volume-only related operations such as create\remove volume.
- The OpenStack OS-Brick component discovers volumes that are attached to hosts, and also controls attach\detach volume operations to\from hosts. When attach\detach volume is being called by either Nova or Cinder, it uses the common code in OS-Brick to perform the action with the ScaleIO driver, which calls add\remove mapped SDC commands.

The OpenStack flow is as follows:

1. Create volume (Cinder operation): Calls the ScaleIO `add_volume` command
2. Attach volume to instance: Initiated by Nova or Cinder:
 - a. The Cinder\Nova driver calls OS-Brick to perform the action.
 - b. The OS-Brick ScaleIO driver detects the volume and then calls the ScaleIO REST command `addMappedSdc`.
3. After the volume is attached to the compute node, both OpenStack and ScaleIO will recognize it as mapped.

Figure 13 OpenStack Liberty integration

ScaleIO with OpenStack Mitaka

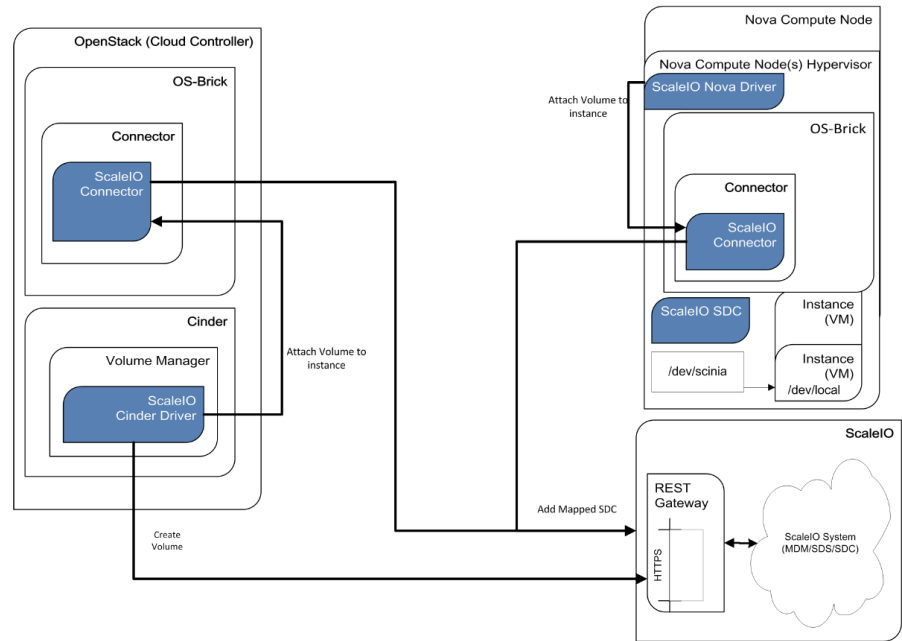
The ScaleIO elastic storage solution for OpenStack Mitaka does not contain an OpenStack installation package, because the ScaleIO drivers are already included in the Mitaka version. For information on configuring the driver, see the OpenStack documentation for the version.

The figure below illustrates the implementation of ScaleIO with OpenStack Mitaka.

- Communication between the Cinder driver and the ScaleIO MDM is implemented using the REST API. This happens usually in volume-only related operations such as create\remove volume.
- The OpenStack OS-Brick component discovers volumes that are attached to hosts, and also controls attach\detach volume operations to\from hosts. When attach\detach volume is being called by either Nova or Cinder, it uses the common code in OS-Brick to perform the action with the ScaleIO driver, which calls add \remove mapped SDC commands.

The OpenStack flow is as follows:

1. Create volume (Cinder operation): Calls the ScaleIO `add_volume` command
2. Attach volume to instance: Initiated by Nova or Cinder:
 - a. The Cinder\Nova driver calls OS-Brick to perform the action.
 - b. The OS-Brick ScaleIO driver detects the volume and then calls the ScaleIO REST command `addMappedSdc`.
3. After the volume is attached to the compute node, both OpenStack and ScaleIO will recognize it as mapped.

Figure 14 OpenStack Mitaka integration

Environment and system requirements

The following environment is supported:

- Liberty and later releases, with either all-in-one OpenStack installation, or multi-node (Compute nodes) installation.
- RHEL 6.x or 7.x (or their CentOS equivalents), or Ubuntu 14.04 or 16.04, using KVM as the hypervisor engine.
- The ScaleIO SDC component must be installed and configured on the OpenStack compute node.
- The ScaleIO gateway (which includes the REST gateway) must be installed as part of the ScaleIO deployment. This gateway is typically installed at the beginning of ScaleIO installation. For more information, see the *ScaleIO Deployment Guide*. The ScaleIO gateway must be accessible from the OpenStack cloud controller. Ensure that you know the IP address of the ScaleIO gateway, and its user name and password, because they are required for the Cinder driver's installation and configuration procedures.
- ScaleIO REST Gateway certificate validation—the OpenStack ScaleIO driver communicates with the REST Gateway through https, (over SSL). By default, the driver ignores verification of the ScaleIO REST Gateway's SSL certificate, but it can verify the certificate if the following configuration parameters are defined:
 - `verify_server_certificate`—set to `True`, if the server's certificate must be verified, and to `False` if no verification is required.
 - `server_certificate_path`—If the parameter `verify_server_certificate` is set to `True`, specify the location of the `.pem` file containing the server's certificate.

For instructions for generating a self-signed certificate using the `keytool` utility, see the *EMC ScaleIO Deployment Guide*.

- The OpenStack driver does not support Storage Pool and domain names that contain special characters, including URL escape characters such as #, /, and so on. If you would like to create a volume on such a Storage Pool or domain in OpenStack, ensure that you specify the Storage Pool or domain using its ID, and not its name.

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."

Configuring Installation Manager properties

Configure the following Installation Manager properties by editing the `gatewayUser.properties` file:

- Enable the Installation Manager (default: true)

To disable, set `features.enable_IM=false`.

You can completely disable the use of the IM's default port, 443, by setting both this property and the `features.enable_gateway` property to false.

- Enable the reuse of previously used devices, and extend a multi-node SDS.

To enable, set `add.sds.with.force.on.next.run=true`.

After finishing and marking the deployment complete, the flag will revert to false, so you will need to set it again, as necessary.

- Exclude RecoverPoint RPA nodes from being upgraded with the Installation Manager.

To exclude RPA nodes while upgrading with the Installation Manager, list the IP addresses to ignore on the `im.ip.ignore.list` line.

Example:

```
im.ip.ignore.list=10.0.0.1,10.0.0.2,...
```

To edit the properties, perform the following:

Procedure

1. Use a text editor to open the `gatewayUser.properties` file, located in the following directory on the Installation Manager/Gateway server:

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

2. Edit the file with the desired changes.
3. Save and close the file.
4. Restart the `scaleio-gateway` service:
 - Windows: Restart the EMC ScaleIO Gateway service
 - Linux: Type the following command:

```
service scaleio-gateway restart
```

Results

Configuration is complete.

CHAPTER 4

Deploying ScaleIO on ESX Servers

This chapter describes how to install, deploy, and perform initial configuration of ScaleIO software components. Topics include:

• Automatic Deployment of ScaleIO on ESXi servers	106
• Deployment prerequisites	106
• Prepare the ESXi environment	109
• Deploying ScaleIO with DirectPath device management	114
• Deploying ScaleIO with RDM/VMDK device management	124
• Install the ScaleIO GUI	134
• Rolling back the deployment wizard in the vSphere Web plug-in	135

Automatic Deployment of ScaleIO on ESXi servers

You can use the ScaleIO vSphere plug-in deployment wizard to deploy ScaleIO on ESXi servers automatically.

ScaleIO provides you with automated tools to install and configure all ScaleIO components on ESXi servers from one workstation, in both physical and virtual environments.

Before you begin, familiarize yourself with the ESX vStorage APIs for Array Integration (VAAI) features supported by the system, as described in the Architecture chapter of your system's Deployment Guide and User Guide.

Note

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 EMC Support at <https://support.emc.com>.

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

Use the ScaleIO vSphere plug-in deployment wizard to install the MDM and SDS components on a dedicated ScaleIO virtual machine (SVM). The SDC is installed directly on the ESX host.

For DirectPath mode deployments, the plug-in can also be used to configure DirectPath on some ESXi servers, and to add SDS devices (after deploying the ScaleIO environment).

The ScaleIO vSphere VMware deployment wizard enables you to perform all of these activities in a simple, efficient manner, over all the machines in a vCenter. The plug-in can also be used for provisioning ScaleIO nodes.

You can also use the deployment wizard with existing ScaleIO systems, to register them to vSphere, and to add SDS server nodes, or extend the existing system.

Note

It is highly recommended to use the deployment wizard for deployment and provisioning. If the need arises, you can perform these tasks manually, or in a combination with the wizard.

Deployment prerequisites

Before you deploy ScaleIO, ensure compliance with the following prerequisites:

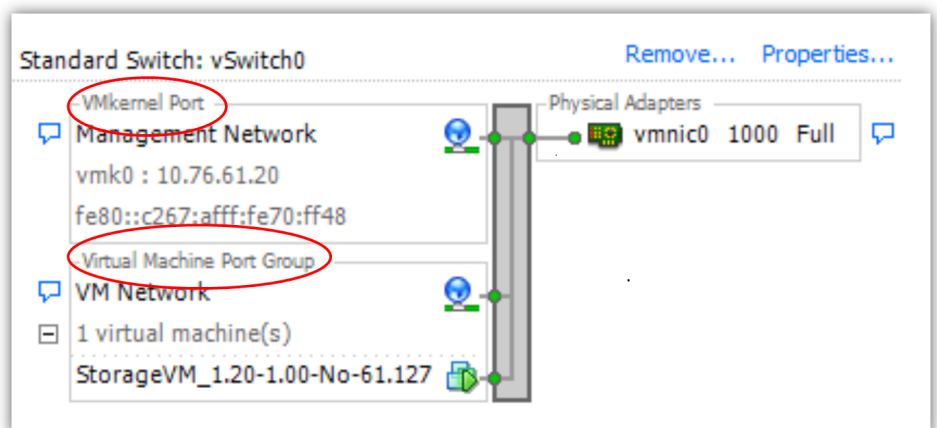
- The ESXi hosts must meet the hardware requirements listed in [System requirements](#) on page 18.
- For deployment in the ScaleIO Ready Node environment, prepare the ESXi hosts before deployment, as described in [Prepare the servers for deployment](#) on page 219.
- All ESXi hosts selected to have either an MDM, Tie-Breaker, or SDS component installed on them, must have a defined local datastore, with a minimum of 10 GB free space (to be used for the SVM). If the ESX is only being used as an SDC, there is no need for this datastore.

- A minimum of three devices to be added to SDS, that all meet the following prerequisites:
 - A minimum of 100 GB available storage capacity.
 - The devices must be free of partitions.
 - If a device is part of a datastore, before adding the device, you must either remove the datastore, or use the plug-in **Advanced settings** option to enable VMDK creation.
 - If the device has the ESXi operating system on it, you must use the plug-in **Advanced settings** option to enable VMDK creation.
- In the following migration or re-installation scenarios, you should contact customer support before deploying:
 - If a datastore was previously used for RDM, and you want to use VMDK devices now.
 - If a datastore was previously used for VMDK, and you want to use RDM devices now.
- The host from which you run the PowerShell (.ps1) script must have the following prerequisites:
 - Runs on Windows, with 64-bit Java installed
 - PowerCLI from VMware (not Windows PowerShell) is installed. PowerCLI version should match the version of the vCenter (for example, both at 6.0).
 - Has incoming and outgoing communication access to the vCenter

Note

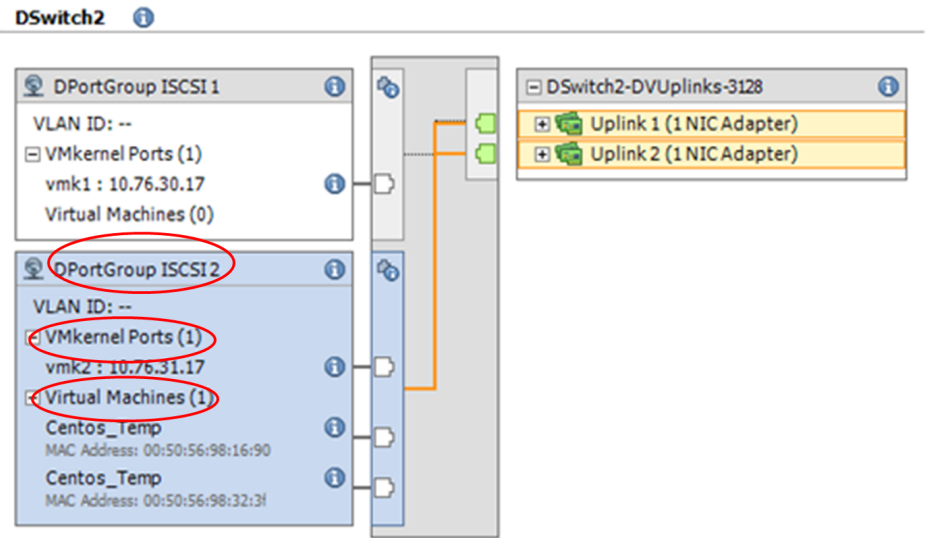
If your vCenter runs on Windows OS, it is recommended to use it.

- The vSphere web client (Virgo) server must have access to the host on which the PowerShell script will be used.
- The management network on all of the ESXs that are part of the ScaleIO system must have the following items configured:
 - Virtual Machine Port Group. (The name must be the same on all of the ESX hosts.)

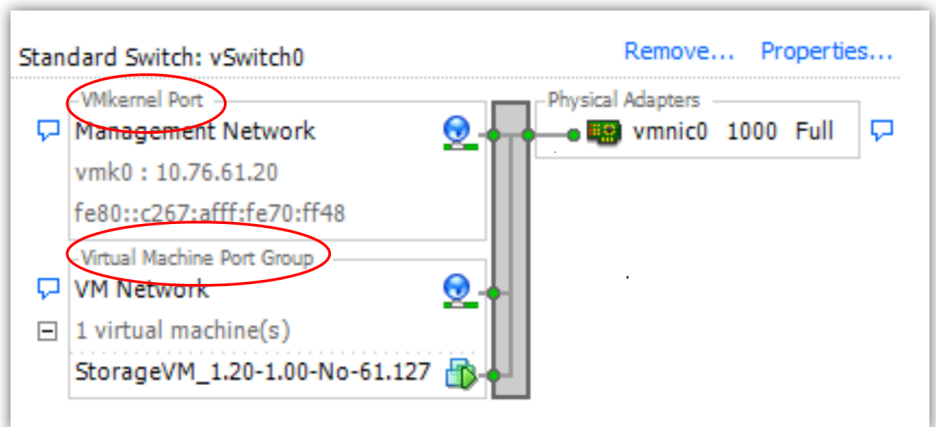


- When using distributed switches, the vDS must have the following items configured:

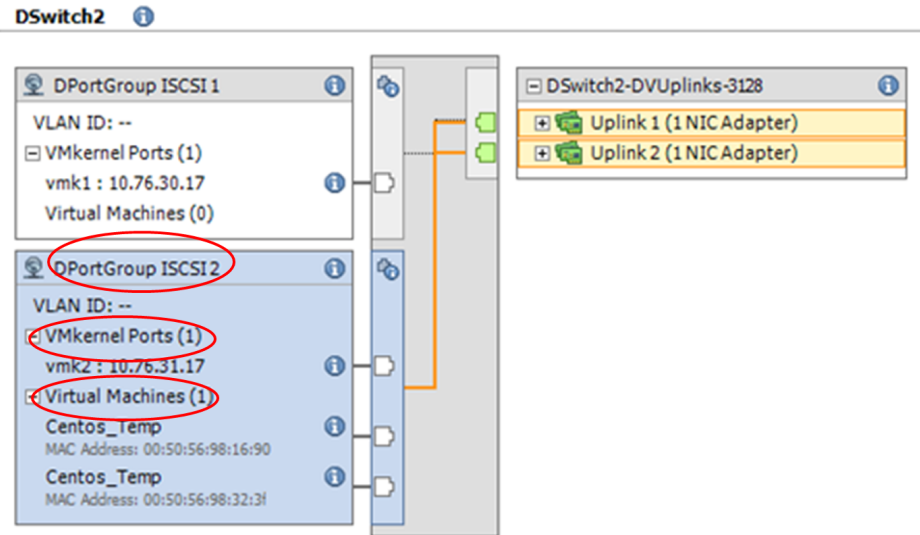
- VMKernel port (necessary only if using a single network)
- dvPortGroup for virtual machines



- Virtual Machine Port Group. (The name must be the same on all of the ESX hosts.)



- When using distributed switches, the vDS must have the following items configured:
 - VMKernel port (necessary only if using a single network)
 - dvPortGroup for virtual machines



- If using only a single network (management), you must manually configure the following:
 - vSwitch
 - VMKernel Port
 - Virtual Machine Port Group
 - VMKernel Port Binding

Prepare the ESXi environment

There are tasks that must be done before deploying ScaleIO in the ESXi environment.

Before deploying ScaleIO in the ESXi environment, you need to prepare the environment:

- Register the ScaleIO vSphere plug-in.
- Upload the OVA (not necessary when deploying on ScaleIO Ready Node servers).
- Access the ScaleIO vSphere plug-in.
- Run the plug-in predeployment steps.

Registering the ScaleIO plug-in

Use PowerCLI to register the ScaleIO plug-in on a vCenter.

Before you begin

Before you begin:

- Ensure that there is communication between the vSphere web client server (usually installed on the vCenter) and the web server storing the plug-in.
- Obtain the vCenter credentials.

To use your own web server, see [“Manual registration of the ScaleIO plug-in”](#).

Procedure

1. Copy the following file to your designated host (preferably your vCenter), that has PowerCLI installed:

EMC-ScaleIO-vSphere-plugin-installer-2.5-<build>.X.zip

2. Extract the contents of the ZIP file.
3. Using PowerCLI for VMware, set to Run as administrator, run the following script:

```
ScaleIOPluginSetup-2.5-<build>.X.ps1
```

- a. Type the vCenter name or IP address, username, and password.

You can also use this script to unregister a plug-in or to create an SVM template.

- b. For Choose mode, select option 1, Register ScaleIO plug-in.

- c. Read the upgrade notice, and type y to continue.

- d. For Select Registration Mode, select Standard (simplified, using embedded Tomcat).

This step may take a few minutes.

Note

You can use the Advanced option to install the plug-in using a ScaleIO Gateway from a previous installation or using your own web service. In either case, you must place this version's plugin.zip file (EMC-ScaleIO-vSphere-web-plugin-2.5-<build>.X.zip) in your resources folder before running the installation. To use a previous ScaleIO Gateway, the resource folder is ScaleIO Gateway installation folder\webapps\root\resources.

- e. If necessary, accept the thumbprint.

4. Log in to the vSphere web client.

If you are already logged in, log out, then log in again.

5. Return to the PowerCLI window, and press ENTER to return to the script menu and complete the plug-in installation.

- 6.



From the vSphere **Home** tab, verify that the ScaleIO icon is visible in the **Inventories** section.

Results

The plug-in is registered. If the ScaleIO icon is missing, the vCenter server failed to register the plug-in, due to one of the following reasons:

- Connectivity problem between the vSphere web client server and the web server storing the plug-in (for example, network / firewall etc.). Resolution: Verify that there is communication between the vSphere web client server and the web server storing the plug-in
- URL problem when using an external web server. Resolution: Verify that the URL is https:// and is pointing to the correct web server IP address (ScaleIO Gateway).

For information on how to use the log to troubleshoot problems that may arise, see [“Troubleshooting plug-in registration issues”](#).

Uploading the OVA template

Use PowerCLI to upload the OVA template to a vCenter. On ScaleIO Ready Node servers, skip this procedure.

Before you begin

Before you begin:

- Ensure that there is communication between the vSphere web client server (usually installed on the vCenter) and the web server storing the plug-in.
- Ensure that the PowerCLI server can run the ScaleIO plugin setup script.
- Obtain the vCenter credentials.

Procedure

1. Copy the following file to your designated host (preferably your vCenter), that has PowerCLI installed:

```
ScaleIOVM_2nics_2.0.14000.X.ova
```

2. Using PowerCLI for VMware, set to Run as administrator, run the following script:

```
ScaleIOPluginSetup-2.5-<build>.X.ps1
```

- a. Type the vCenter name or IP address, user name, and password.
- b. From the PowerCLI script, select 3 to Create SVM template.
- c. Type the parameters described in the following table. (If you are already logged in, some of these parameters may not be necessary.)

Parameter	Description
vcenter	vCenter name or IP address
user name	vCenter user name
password	vCenter password
datacenter	The name of the data center where the datastore that will store the template resides
ova_path	The path of the SVM's OVA
datastores	A list of datastores, up to eight, on which the templates will be created. Press ENTER to stop specifying datastores.

Note

For best results, enter a local (not shared) datastore for each ESX server.

For faster, parallel, deployment in large-scale environments, you can use the OVA to create SVM templates on as many as eight datastores. To do so,

type the datastore names, and when you are done, leave the next line blank. The following example shows how to enter two datastores:

```
datastores[0]: datastore1
datastores[1]: datastore1 (1)
datastores[2]:
```

The upload procedure can take several minutes, during which time a temporary SVM is created, the templates are created, then the temporary SVM is deleted.

When each template is created, a message similar to the following appears:

```
The template EMC ScaleIO SVM Template (v2.0.14000.X) was
successfully created
```

d. When the process is complete, type 4 to exit the plug-in script.

Accessing the plug-in

Access the ScaleIO plug-in from the vSphere Web Client home screen.



From the vSphere Web Client home tab, click

The **EMC ScaleIO** screen appears:

Figure 15 EMC ScaleIO screen



Preparing the ESXi hosts

Use the vSphere plug-in to prepare the ESXi hosts for ScaleIO deployment.

This procedure is a prerequisite for deploying a new ScaleIO system, and in some cases for extending an existing system.

Preparing the ESXi hosts includes the following:

- **Install SDC on the host.**
This is required for every ESX host that will be included in the system, and it is strongly recommended to enable this for every host that might be part of the system in the future.
- **Configure DirectPath**
Enables the SVM to take control of ESX devices. If the host has multiple controllers, you must configure DirectPath on that host manually, not with the plug-in. After that is done, you can use the plug-in to deploy ScaleIO on that host and to add devices to it.

When DirectPath is configured, all servers in the system must use DirectPath. For ScaleIO Ready Node or VxRack Node 100 Series systems, this can be configured with the ScaleIO vSphere plug-in. For other servers, use the vSphere client to configure each ESX host manually.
- **Convert ISO**
ScaleIO Ready Node 13G and 14G nodes come with a pre-installed ISO. To add these to a system with the vSphere plug-in, you must select the **Convert ISO** option for each node.

Procedure

1. From the **Basic tasks** section of the **EMC ScaleIO** screen, click **Pre-Deployment Actions**.

The **Pre-Deployment Actions** screen appears, showing all the ESX hosts on all the connected vCenters.

2. Select the ESX hosts, and select the settings required for each.

Note

For the **Install SDC** option, it is highly recommended to select all ESX hosts that may be included in an ESX system, even if only in the future.

3. For ScaleIO Ready Node 13G and 14G nodes, you must select **Convert ISO** for each node.
4. Type the root password for each host.

All changes or entries made at the vCenter level will be applied to all servers in the cluster.
5. Click **Run**.

The status appears in the dialog.
6. When finished, click **Finish**.
7. Restart each ESX host.

Note

You must restart the ESX hosts before proceeding.

After rebooting, a RAID controller that was configured with DirectPath will be displayed in the vSphere client **Configure** tab, on the **DirectPath I/O PCI Devices Available to VMs** screen.

8. When using all-SSD chassis, perform this step.

Sometimes the ESX nodes in an all-SSD chassis do not identify the node disks as Standard Parallel SCSI devices. If so, you must enable the **RDMs on non Parallel SCSI Controllers** option.

- a. In the **EMC ScaleIO** screen, click **Advanced Settings** to display the settings options.
- b. Select the **Enable RDMs on non Parallel SCSI Controller** option and click **OK**.

Results

After finishing this task, the results of your selections are displayed after reopening the **Pre-Deployment Actions** screen.

After you finish

Proceed with the ScaleIO deployment.

Deploying ScaleIO with DirectPath device management

Use the VMware deployment wizard to deploy ScaleIO when devices are managed by VMware DirectPath I/O.

Before you begin

- Ensure that all ESXi servers to be added to the system were pre-configured to enable DirectPath.
- Configure advanced installation options (optional):
 - Allow taking over devices that were used in other ScaleIO systems.
 - Allow the use of non-local datastores for the Gateway.
 - Increase parallelism limit.

To access these settings, click **Advanced settings** on the **EMC ScaleIO** screen. For more information, see [“Advanced settings options”](#).

Procedure

1. From the **Basic tasks** section of the screen, click **Deploy ScaleIO environment**.

The ScaleIO VMware deployment wizard begins. If you exited the previous deployment before completion, you will be able to return from where you left off.

NOTICE

The deployment wizard assumes that you are using the provided ScaleIO OVA template to create the ScaleIO virtual machines.

2. In the **Select Installation** screen, start the deployment of a new system:
 - a. Select **Create new ScaleIO system**.
 - b. Review and approve the license terms.
 - c. Click **Next**.
3. In the **Create New System** screen, type the following, then click **Next**:
 - **System Name:** Type a unique name for this system.
 - **Admin Password:** Type and confirm a password for the ScaleIO admin user. The password must meet the listed criteria.
4. In the **Add ESX Hosts to Cluster** screen, select the ESX hosts to add as part of the system:
 - a. Select the vCenter on which to deploy the ScaleIO system.
The vCenter information is populated in the lower part of the screen.
 - b. Select **Set up a DirectPath-based system**.
 - c. Expand the vCenter, select the ESX hosts to add to the ScaleIO system, then click **Next**.

Note

To configure ScaleIO, you must select a minimum of three ESX hosts. ESX hosts that do not have the SDC installed, or hosts for which DirectPath was configured before deployment, but DirectPath was not selected in the previous step will not be available.

The **Select Management Components** screen appears:

Select ESX hosts for ScaleIO components:

☒ 3-node mode
☐ 5-node mode

Initial Master MDM:

Manager MDM:

TieBreaker MDM:

Optional:

Standby Manager:

Standby Manager:

Standby TB:

Standby TB:

5. Configure the management components:
 - a. Select to deploy a 3-node or 5-node cluster.
The next fields on this screen will change, depending on your choice.

- b. Select an ESX server to serve for each of the MDM cluster roles.
You can give a name to the MDM servers, such as Manager1, etc.
- c. Select ESX servers to serve as Standby manager and tiebreaker roles (optional).
- d. Click **Next**.

The **Configure Performance, Sizing, and Syslog** screen appears.

6. Configure the following settings (optional), then click **Next**:
 - To configure components for high-performance, select them.
 - To configure the allocation of SVM RAM, select from the following:
 - To use default RAM allocation, select **Standard size**.
 - To use custom settings, select **Custom size**, and type the maximum capacity and maximum number of volumes.
 - To configure syslog reporting, select **Configure syslog**, and type the syslog server, port (default: 1468), and facility (default: 0).
 - To configure DNS servers, type their details.

The **Configure Protection Domains** screen appears:

Add new Protection Domain to ScaleIO system:

Protection Domain name:

RAM Read Cache size per SDS:

128 MB

▲▼

Add

Existing Protection Domains in ScaleIO system:

Protection Domain Name	RAM Read Cache Size (in MB)
Deploy_assist_01	128

Remove

You can create (or remove) Protection Domains (PD). You must create at least one PD.

7. Create a Protection Domain:

a. Enter the following information:

- **Protection Domain name:** It is recommended to use a meaningful name.
- **RAM Read Cache size per SDS:** Minimum 128 MB (You can increase this for your environment needs.)

b. Click **Add**.

The added PDs appear in the lower section of the screen, together with the existing PDs. To remove a newly created PD, select it and click **Remove**.

c. To create additional PD, repeat this step.

d. Click **Next**.

The **Configure Storage Pools** screen appears.

In the **Configure Storage Pools** screen, you can create (or remove) Storage Pools (SP). You must create at least one SP.

8. Create a Storage Pool:

a. Type the **Storage Pool name:** It is recommended to use meaningful names.

b. Select to which PD to add the SP.

c. To enable zero padding, select **Enable zero padding**. Zero padding must be enabled for use with RecoverPoint replication and for using the background scanner in data comparison mode.

d. To enable Read Flash cache, select **Enable RFlash**.

e. Click **Add**.

The added SPs appear in the lower section of the screen, together with the existing PDs. To remove a newly created SP, select it and click **Remove**.

f. To create additional SPs, repeat this step.

g. Click **Next**.

The **Create Fault Sets** screen appears. You can use this screen to create Fault Sets (optional).

Note

When defining Fault Sets, you must follow the guidelines described in [“Fault Sets”](#). Failure to do so may prevent creation of volumes.

9. Create a Fault Set (optional):

a. Type the **Fault Set name**. It is recommended to use meaningful names.

b. Select to which PD to add the Fault Set.

c. Click **Add**

Added Fault Sets appear in the lower section of the screen, inside the folder of the parent PD. You can remove a newly created Fault Set by selecting it and clicking **Remove**.

d. Repeat to create additional Fault Sets (minimum of three), then click **Next**.

The **Add SDSs** screen appears.

10. Configure the following for every ESX host or SVM, then click **Next**:
 - a. For every SVM in a DirectPath deployment, you must select **SDS** and assign a **Protection Domain**.

Note

To make the same selections for every ESX in a cluster, you can make your selections per cluster or datacenter.

- b. You can select a **Fault Set** (optional).
 - c. Type the **Total Planned Capacity** for the SDS, even if it is only being partially populated now. This sets the memory allocation for the SVM, thus avoiding manual memory calculations.
- If you don't configure this now, you can set the SVM memory later manually, as described in "SVM manual memory allocation."

- d. Click **Next**.

Adding devices to SDS is done after the deployment is complete.

The **Add SDCs** screen appears.

11. Configure the SDCs:
 - a. For each ESX to be added as an SDC:
 - a. Select the **SDC** check box.
 - b. Type the ESX root password.

Note

To show the entered ESX passwords, select **Show passwords**.

- b. Choose whether to enable or disable the LUN comparison for ESX hosts.
- In general, in environments where the SDC is installed on ESX and also on physical hosts, you should set this to **Disable**.

Note

Before enabling LUN comparison, consult your environment administrator.

- c. Click **Next**.

The **Configure Upgrade Components** dialog appears.

12. Configure the ScaleIO Gateway and LIA:
 - a. Select an ESX to host the Gateway virtual machine.

A unique SVM will be created for the Gateway.
 If the previously-selected ESX servers do not have sufficient free space (on any datastore) to contain the ScaleIO SVM template, an SVM, and the Gateway SVM, you will not have an option to select an ESX in this step. It will be done automatically.

- b. Type and confirm a password for the Gateway administrative user.
- c. Type and confirm a password for the LIA.
The password must be the same across all SVMs in the system.
- d. Click **Next**.

Note

You can only move forward if the passwords meet the listed criteria, and if the confirmation passwords match the entered passwords.

The **Select OVA Template** screen appears:

13. Configure templates:

- a. Select the template to use to create the ScaleIO virtual machines (SVM).

The default is **EMC ScaleIOSVM Template**. If you uploaded a template to multiple datastores, you can select them all, for faster deployment.

If the Gateway selection was performed automatically in the previous step (indicating insufficient space), you must choose at least two templates in this step, one of which will be converted to the Gateway SVM.

After selecting the templates, the deployment wizard will automatically select one of the ESXs with the templates to host the Gateway and during deployment will convert the template to a VM for the Gateway (instead of cloning the template).

Note

If you select a custom template, ensure that it is compatible with the VMware plug-in and the ScaleIO MDM.

- b. Type and confirm a new root password that will be used for all SVMs to be created.
- c. Click **Next**.

The **Configure Networks** screen appears:

Select networks that will be used on ScaleIO VMs:

IPv4	Management network label:	VM Network	
IPv4	Data network label:	data	Create new network
	2nd data network label:	N/A	Create new network

*Note that "Create new network" dialog will create vSwitches (not distributed switches)

14. Select the network configuration. You can select an existing (simple or distributed) network, or select **Create a new network**.

You can use a single network for management and data transfer or separate networks. Separating the networks is recommended for security and increased efficiency. You can select one data network or two.

The management network, used to connect and manage the SVMs, is normally connected to the client management network, a 1 GB network.

The data network is internal, enabling communication between the ScaleIO components, and is recommended to be at least a 10GB network.

Note

The selected networks must have communication with all of the system nodes. In some cases, while the wizard does verify that the network names match, this does not guarantee communication, as the VLAN IDs may have been manually altered.

- a. To use one network, select a protocol (IPv4 or IPv6) and a management network, click **Next** and proceed with SVM configuration.

NOTICE

When the MDM and SDS are on physical servers, IPv6 is fully supported. Manually configure the IPv6 on the ESX (VMKernel). IPv6 is not supported in VMware converged solutions. Contact EMC support for more information.

For best results, it is highly recommended to use the plug-in to create the data networks, as opposed to creating them manually.

- b. To use separate networks, select a protocol (IPv4 or IPv6) for the management network label and one or two data network labels. If the data network already exists (such as a customer pre-configured distributed switch or a simple vswitch), select it from the drop-down box. Otherwise, configure the data network by clicking **Create new network**.

The **Create New Data Network** screen appears:

Create New Data Network

Network name:

VMkernel name:

VLAN ID:

Network type: IPv4 IPv6

ESX Name		VMkernel IP	VMkernel Subnet Mask
10.76.61.8	<input type="text"/>	<input type="text"/>	<input type="text"/>
10.76.61.9	<input type="text"/>	<input type="text"/>	<input type="text"/>
10.76.61.32	<input type="text"/>	<input type="text"/>	<input type="text"/>

*Note that vSwitches (not distributed switches) will be created

OK Cancel

c. Configure the networks:

- **Network name:** The name of the VMware network
- **VMkernel name:** The name of the VMkernel (used to support multipathing)
- **VLAN ID:** the network ID
- **Network type:** IPv4 or IPv6
- For each ESX, select a **Data NIC**, a **VMkernel IP**, and a **VmKernel Subnet Mask**.

Note

You can click  to auto-fill the values for **Data NIC** and **VMkernel IP**.

d. Click **OK**.

The data network is created.

The wizard will automatically configure the following for the data network:

- vSwitch
- VMkernel Port
- Virtual Machine Port Group
- VMkernel Port Binding

e. Click **Next**.

The **Configure SVM** screen appears.

ScaleIO VMware Installation Wizard

1. Select installation type
2. Confirm license
3. Create new system
4. Add ESXi hosts to cluster
5. Select management components
6. Configure Performance, Storage, System
7. Configure Protection Domains
8. Configure Storage Pools
9. Create Fault Sets (optional)
10. Add SDCs
11. Add devices to SDCs
12. Add SDCs
13. Configure Upgrade Components
14. Select OVA template
15. Configure networks
16. Configure SVM

Configure ScaleIO Virtual Machine IP addresses and Hosting Datastore

ESXi Name	Management IP	Management Subnet Mask	Default Gateway	Data IP	Data Subnet Mask
10.100.150.133 (ScaleIO Gateway)					
10.100.150.71 (Master MDM)					
10.100.150.70 (Slave 1 MDM)					
10.100.150.133 (WebBreaker 1)					

Configure Cluster Virtual IP addresses

Network	Virtual IP
Data (data0)	
2nd Data (data02)	

Back Next Finish Cancel

15. Configure all the SVMs:

- Type the IP address, Subnet mask, and Default Gateway for the management network, then the data network.
- Type the Cluster Virtual IP for each network interface.
- You can select a datastore, or allow automatic selection.
- Configure the cluster's virtual IP addresses by entering the virtual IP address for each data network.
- Click **Next**.

Icons indicate the role that the server plays in the ScaleIO system. You can

select  to auto-fill the values for IP.

The **Review Summary** screen appears.

16. Review the configuration.

Click **Finish** to begin deployment or **Back** to make changes.

17. Type the vCenter user name and password, then click **OK** to begin the deployment.

The **Deployment Progress** screen appears.

During the deployment process you can view progress, pause the deployment, and view logs.

To pause the deployment, click **Pause**. Steps that are already in progress will pause after they complete.

After pausing, select one of the following options:

- **Continue deployment** to continue.
- **Abort** to abort the deployment process.
- **Cancel and Rollback entire deployment** to roll back all deployment activities.

Note

Rollback cannot be canceled once started.

For more details on rolling back deployment, see [Rolling back the deployment wizard in the vSphere Web plug-in](#) on page 135

- **Rollback failed tasks** to roll back only the tasks that failed.
-

Note

Rollback cannot be canceled once started.

18. When the deployment is complete, click **Finish**.

If a task failed, click **Continue deployment** to try again.

After you finish

After deployment is complete, set all SVMs to start automatically with the system. Do not set SVMs under the VMware resource-pool feature.

Use the vSphere plug-in to add devices to SDS.

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.

Deploying ScaleIO with RDM/VMDK device management

Use the VMware deployment wizard to deploy ScaleIO when devices are configured as RDM or VMDK.

Before you begin

Configure advanced installation options (optional):

- Ensure that all ESXi servers to be added to the system were pre-configured with the pre-deployment steps.
- Configure advanced installation options (optional):
 - Enable creation of VMDK.
 - Enable RDMs on non-parallel SCSI controllers.
 - Allow taking over devices that were used in other ScaleIO systems.
 - Allow the use of non-local datastores for the Gateway.
 - Increase parallelism limit.

To access these settings, click **Advanced settings** on the **EMC ScaleIO** screen. For more information, see [“Advanced settings options”](#).

To deploy in a 2-layer environment, where only the SDCs are on ESXi servers, see [Deploying ScaleIO in a 2-layer environment using the Installation Manager and plug-in](#) on page 217.

Procedure

1. From the **Basic tasks** section of the screen, click **Deploy ScaleIO environment**.

The ScaleIO VMware deployment wizard begins. If you exited the previous deployment before completion, you will be able to return from where you left off.

NOTICE

The deployment wizard assumes that you are using the provided ScaleIO OVA template to create the ScaleIO virtual machines.

2. In the **Select Installation** screen, start the deployment of a new system:
 - a. Select **Create new ScaleIO system**.
 - b. Review and approve the license terms.
 - c. Click **Next**.
3. In the **Create New System** screen, type the following, then click **Next**:
 - **System Name:** Type a unique name for this system.
 - **Admin Password:** Type and confirm a password for the ScaleIO admin user. The password must meet the listed criteria.
4. In the **Add ESX Hosts to Cluster** screen, select the ESX hosts to add as part of the system:
 - a. Select the vCenter on which to deploy the ScaleIO system.
The vCenter information is populated in the lower part of the screen.
 - b. Expand the vCenter, select the ESX hosts to add to the ScaleIO system, then click **Next**.

Note

To configure ScaleIO, you must select a minimum of three ESX hosts. ESX hosts that do not have the SDC installed, or hosts for which DirectPath was configured before deployment, but DirectPath was not selected in the previous step will not be available.

The **Select Management Components** screen appears:

Select ESX hosts for ScaleIO components:

☒ 3-node mode
☐ 5-node mode

Initial Master MDM:	10.108.159.70	Manager1
Manager MDM:	10.108.159.71	Manager2
TieBreaker MDM:	10.108.158.133	TB1

Optional:

Standby Manager:		Standby1
Standby Manager:		Standby2
Standby TB:		Standby3
Standby TB:		Standby4

5. Configure the management components:

- a. Select to deploy a 3-node or 5-node cluster.

The next fields on this screen will change, depending on your choice.

- b. Select an ESX server to serve for each of the MDM cluster roles.

You can give a name to the MDM servers, such as Manager1, etc.

- c. Select ESX servers to serve as Standby manager and Tie Breaker roles (optional).

- d. Click
- Next**
- .

The **Configure Performance, Sizing, and Syslog** screen appears.

6. Configure the following settings (optional), then click **Next**:

- To configure components for high-performance, select them.
- To configure the allocation of SVM RAM, select from the following:
 - To use default RAM allocation, select **Standard size**.
 - To use custom settings, select **Custom size**, and type the maximum capacity and maximum number of volumes.
- To configure syslog reporting, select **Configure syslog**, and type the syslog server, port (default: 1468), and facility (default: 0).
- To configure DNS servers, type their details.

The **Configure Protection Domains** screen appears:

Add new Protection Domain to ScaleIO system:

Protection Domain name:

RAM Read Cache size per SDS:

128 MB

Add

Existing Protection Domains in ScaleIO system:

Protection Domain Name	RAM Read Cache Size (in MB)
Deploy_assist_01	128

Remove

You can create (or remove) Protection Domains (PD). You must create at least one PD.

7. Create a Protection Domain:

a. Enter the following information:

- **Protection Domain name:** It is recommended to use a meaningful name.
- **RAM Read Cache size per SDS:** Minimum 128 MB (You can increase this for your environment needs.)

b. Click **Add**.

The added PDs appear in the lower section of the screen, together with the existing PDs. To remove a newly created PD, select it and click **Remove**.

c. To create additional PD, repeat this step.

d. Click **Next**.

The **Configure Storage Pools** screen appears.

In the **Configure Storage Pools** screen, you can create (or remove) Storage Pools (SP). You must create at least one SP.

8. Create a Storage Pool:

a. Type the **Storage Pool name:** It is recommended to use meaningful names.

b. Select to which PD to add the SP.

c. To enable zero padding, select **Enable zero padding**. Zero padding must be enabled for use with RecoverPoint replication and for using the background scanner in data comparison mode.

d. To enable Read Flash cache, select **Enable RFlash**.

e. Click **Add**.

The added SPs appear in the lower section of the screen, together with the existing PDs. To remove a newly created SP, select it and click **Remove**.

f. To create additional SPs, repeat this step.

g. Click **Next**.

The **Create Fault Sets** screen appears. You can use this screen to create Fault Sets (optional).

Note

When defining Fault Sets, you must follow the guidelines described in [“Fault Sets”](#). Failure to do so may prevent creation of volumes.

9. Create a Fault Set (optional):

a. Type the **Fault Set name**. It is recommended to use meaningful names.

b. Select to which PD to add the Fault Set.

c. Click **Add**

Added Fault Sets appear in the lower section of the screen, inside the folder of the parent PD. You can remove a newly created Fault Set by selecting it and clicking **Remove**.

d. Repeat to create additional Fault Sets (minimum of three), then click **Next**.

The **Add SDSs** screen appears.

10. Configure the following for every ESX host or SVM:

- a. Select **SDS** to assign an SDS role.

Note

To make the same selections for every ESX in a cluster, you can make your selections per cluster or datacenter.

- b. If the node is an SDS, assign a **Protection Domain**.

- c. You can select a **Fault Set** (optional).

- d. Click **Next**.

The **Add devices to SDSs** screen appears, showing the clusters.

This screen has the following tabs:

- **Information** - shows the selected ESX and cluster.
- **Assign devices** - select hosts and assign devices.
- **Replicate selections** - replicate device selections from one host to others. This can be very useful if your ESXs have identical attached devices. For example, if you select an SSD device for the source ESX, and then replicate this selection to the targets, the deployment wizard can automatically select all other SSD devices on the target SDSs.

Device matching is performed based on the device runtime name.

To replicate device selections, all of the following conditions must be met:

- The number of devices on each ESX must be identical.
- Source and target devices must be identical in the following ways: a) both are SSD or non-SSD, b) both have datastores on them or do not, c) both are roughly the same size (within 20%), and d) both are connected via a RAID controller or directly attached.
- At least one of the following conditions must be met: a) both SDSs are in the same Protection Domain, b) both SDSs are in different Protection Domains, but with the same list of Storage Pools, or c) the target SDS is in a Protection Domain with only one Storage Pool.

11. On the **Information** tab, select an ESX host from a cluster, then click **Assign devices**.

The **Assign devices** tab appears.

This screen shows the devices whose free space can be added to the selected ESX host/SDS. You should balance the capacity over the selected SDS.

12. To assign a device's space to an SDS, perform the following:

- a. In the **Use for** drop-down, select **Storage** or **RFcache**.
- b. In the **Pool Name** drop-down, select the Storage Pool (SP) to which to assign the device.

Note

If the selected SP has RFlcache enabled, you must select at least one RFlcache device for that SDS node.

TIP: You can select all available devices by clicking **Select all devices**, and selecting their use and Storage Pool.

Note

If you selected to create VMDK (before the deployment), the following options appear:

- **Create VMDK.** Select this for all relevant devices.
 - **Select all available devices.** Click this to select all devices with a VMFS, and with unused capacity that can be added to the ScaleIO system.
-

- c. If the server is being only partially populated now, you can prevent future server restarts by allocating additional RAM during the deployment.
- d. Click **Assign**.
13. To replicate selections to other SDSs, perform the following:
 - a. Select the **Replicate selection** tab.
 - b. Select the ESX whose device selection you wish to replicate.
This is the source ESX.
 - c. Select the target ESXs to which to replicate the selection of the source ESX.
 - d. Click **Copy configuration**.
The results are displayed in the right pane of the screen.
14. When you have selected devices for all SDSs, click **Next**.

Note

You must select at least one device for each SDS.

The **Add SDCs** screen appears.

15. Configure the SDCs:
 - a. For each ESX to be added as an SDC:
 - a. Select the **SDC** check box.
 - b. Type the ESX root password.
-

Note

To show the entered ESX passwords, select **Show passwords**.

- b. Choose whether to enable or disable the LUN comparison for ESX hosts.
In general, in environments where the SDC is installed on ESX and also on physical hosts, you should set this to **Disable**.

Note

Before enabling LUN comparison, consult your environment administrator.

- c. Click **Next**.

The **Configure Upgrade Components** dialog appears.

16. Configure the ScaleIO Gateway and LIA:

- a. Select an ESX to host the Gateway virtual machine.

A unique SVM will be created for the Gateway.

If the previously-selected ESX servers do not have sufficient free space (on any datastore) to contain the ScaleIO SVM template, an SVM, and the Gateway SVM, you will not have an option to select an ESX in this step. It will be done automatically.

- b. Type and confirm a password for the Gateway administrative user.

- c. Type and confirm a password for the LIA.

The password must be the same across all SVMs in the system.

- d. Click **Next**.
-

Note

You can only move forward if the passwords meet the listed criteria, and if the confirmation passwords match the entered passwords.

The **Select OVA Template** screen appears:

17. Configure templates:

- a. Select the template to use to create the ScaleIO virtual machines (SVM).

The default is **EMC ScaleIOSVM Template**. If you uploaded a template to multiple datastores, you can select them all, for faster deployment.

If the Gateway selection was performed automatically in the previous step (indicating insufficient space), you must choose at least two templates in this step, one of which will be converted to the Gateway SVM.

After selecting the templates, the deployment wizard will automatically select one of the ESXs with the templates to host the Gateway and during deployment will convert the template to a VM for the Gateway (instead of cloning the template).

Note

If you select a custom template, ensure that it is compatible with the VMware plug-in and the ScaleIO MDM.

- b. Type and confirm a new root password that will be used for all SVMs to be created.

- c. Click **Next**.

The **Configure Networks** screen appears:

Select networks that will be used on ScaleIO VMs:

IPv4	Management network label:	VMNetwork	
IPv4	Data network label:	data	Create new network
	2nd data network label:	N/A	Create new network

**Note that "Create new network" dialog will create vSwitches (not distributed switches)*

18. Select the network configuration. You can select an existing (simple or distributed) network, or select **Create a new network**.

You can use a single network for management and data transfer or separate networks. Separating the networks is recommended for security and increased efficiency. You can select one data network or two.

The management network, used to connect and manage the SVMs, is normally connected to the client management network, a 1 GB network.

The data network is internal, enabling communication between the ScaleIO components, and is recommended to be at least a 10GB network.

Note

The selected networks must have communication with all of the system nodes. In some cases, while the wizard does verify that the network names match, this does not guarantee communication, as the VLAN IDs may have been manually altered.

- a. To use one network, select a protocol (IPv4 or IPv6), and a management network, then proceed to the next step, configuring the SVMs.

NOTICE

When the MDM and SDS are on physical servers, IPv6 is fully supported. Manually configure the IPv6 on the ESX (VMKernel). IPv6 is not supported in VMware converged solutions. Contact EMC support for more information.

For best results, it is highly recommended to use the plug-in to create the data networks, as opposed to creating them manually.

- b. To use separate networks, select a protocol (IPv4 or IPv6) for the management network label and one or two data network labels. If the data network already exists (such as a customer pre-configured distributed switch or a simple vswitch), select it from the drop-down box. Otherwise, configure the data network by clicking **Create new network**.

The **Create New Data Network** screen appears:

Network name:

VMkernel name:

VLAN ID:

Network type: IPv4 (dropdown menu open showing IPv4 and IPv6)

ESX Name	Data NIC	VMkernel IP	VMkernel Subnet Mask
10.76.61.8	<input type="text"/>	<input type="text"/>	<input type="text"/>
10.76.61.9	<input type="text"/>	<input type="text"/>	<input type="text"/>
10.76.61.32	<input type="text"/>	<input type="text"/>	<input type="text"/>

**Note that vSwitches (not distributed switches) will be created*

OK Cancel

c. Configure the networks:

- **Network name:** The name of the VMware network
- **VMkernel name:** The name of the VMkernel (used to support multipathing)
- **VLAN ID:** the network ID
- **Network type:** IPv4 or IPv6
- For each ESX, select a **Data NIC**, a **VMkernel IP**, and a **VMkernel Subnet Mask**.

Note

You can click  to auto-fill the values for **Data NIC** and **VMkernel IP**.

d. Click **OK**.

The data network is created.

The wizard will automatically configure the following for the data network:

- vSwitch
- VMkernel Port
- Virtual Machine Port Group
- VMkernel Port Binding

e. Click **Next**.

The **Configure SVM** screen appears.

ScaleIO VMWare Installation Wizard

1. Select installation type
2. Confirm license
3. Create new system
4. Add SDC hosts to cluster
5. Select management components
6. Configure Performance, Billing, Storage
7. Configure Protection Domains
8. Configure Storage Pools
9. Create Fault Sets (optional)
10. Add SDCs
11. Add devices to SDCs
12. Add SDCs
13. Configure Upgrade Components
14. Select OVA template
15. Configure networks
16. Configure SVM

Configure ScaleIO Virtual Machine IP addresses and Hosting Datastore

ESX Name	Management IP	Management Subnet Mask	Default Gateway	Data IP	Data Subnet Mask
10.100.150.133 (ScaleIO Gateway)					
10.100.150.71 (Master MDM)					
10.100.150.70 (Slave 1 MDM)					
10.100.150.133 (WebBreaker 1)					

Configure Cluster Virtual IP addresses

Network	Virtual IP
Data (data0)	
2nd Data (data02)	

Back Next Finish Cancel

19. Configure all the SVMs:

- Type the IP address, Subnet mask, and Default Gateway for the management network, then the data network.
- Type the Cluster Virtual IP for each network interface.
- You can select a datastore, or allow automatic selection.
- Configure the cluster's virtual IP addresses by entering the virtual IP address for each data network.
- Click **Next**.

Icons indicate the role that the server plays in the ScaleIO system. You can

select  to auto-fill the values for IP.

The **Review Summary** screen appears.

20. Review the configuration.

Click **Finish** to begin deployment or **Back** to make changes.

21. Type the vCenter user name and password, then click **OK** to begin the deployment.

The **Deployment Progress** screen appears.

During the deployment process you can view progress, pause the deployment, and view logs.

To pause the deployment, click **Pause**. Steps that are already in progress will pause after they complete.

After pausing, select one of the following options:

- **Continue deployment** to continue.
- **Abort** to abort the deployment process.
- **Cancel and Rollback entire deployment** to roll back all deployment activities.

Note

Rollback cannot be canceled once started.

For more details on rolling back deployment, see [Rolling back the deployment wizard in the vSphere Web plug-in](#) on page 135

- **Rollback failed tasks** to roll back only the tasks that failed.
-

Note

Rollback cannot be canceled once started.

22. When the deployment is complete, click **Finish**.

If a task failed, click **Continue deployment** to try again.

After you finish

Note

After deployment is complete, set all SVMs to start automatically with the system. Do not set SVMs under the VMware resource-pool feature.

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."

Rolling back the deployment wizard in the vSphere Web plug-in

When deployment fails due to incorrect input or failed devices, you can roll back the deployment wizard, fix or change the specific deployment setting, and restart deployment.

Procedure

1. When deployment fails, click **Cancel and roll back entire deployment**.
2. In the **Rollback settings** dialog box, select whether you want to remove the SDC configuration.
 - If you select **Remove SDCs configuration (requires a manual restart of the ESXs)**, the ScaleIO VMs are deleted, the MDM IPs are removed from the ESXs, and you must manually restart the ESXs.
 - If you do not select **Remove SDCs configuration (requires a manual restart of the ESXs)**, the ScaleIO VMs are deleted, the MDM IPs are not removed from the ESXs, and you do not need to manually restart the ESXs.
3. When the deployment roll back is complete, select **Finish deployment and return to wizard**.

If you select **Finish**, you will not be returned to the wizard and will have to start the deployment from the beginning; your deployment details will not be saved. After confirming that you want to finish the deployment and return to the wizard, the last screen in the deployment wizard is displayed.

4. Click **Back** or **Next** to navigate backwards or forwards through the deployment wizard.
Change the settings that caused the initial deployment to fail.
5. Click **Finish** to complete the wizard and begin deployment.

CHAPTER 5

Post-Deployment Activities

This chapter describes activities that are performed after deployment of a ScaleIO system. Topics include:

• Post-deployment checklist	138
• Create a Lockbox	139
• Configuring ESRS	140
• Configure native users	141
• Configure LDAP users	142
• Configuring SNMP after deployment	143
• Create and map volumes	144

Post-deployment checklist

Use the following checklist to verify that you complete the required and optional post-deployment procedures.

Post-deployment procedures

Procedure	Notes	Mandatory?	Done
Install ScaleIO license	A ScaleIO license is necessary when deploying ScaleIO in a production environment.	Optional	
Review ScaleIO performance best practices	See the <i>ScaleIO Performance Fine-Tuning Technical Notes</i> .	Recommended	
Review ScaleIO security best practice configurations	See the <i>ScaleIO Security Configuration Guide</i> .	Recommended	
Back up all SSH keys from nodes and ScaleIO VMs		Optional	
Add SDS devices	After deploying ScaleIO on ESXi servers with DirectPath device management, use the vSphere plug-in to devices to the SDS.	Mandatory	
Run the system analysis tool from the Installation Manager (IM), prior to provisioning.	It is highly recommended to run the system analysis tool immediately after deployment, before provisioning volumes, and before using the system in production. You can also use it to check the health of a system that is already operational. The system analysis tool supports only RHEL 6.x servers and IPv4 network configuration. For more information, see System analysis overview on page 240.	Highly recommended	
Create a LockBox and add the MDM credentials to it	A LockBox can be created during installation with the Installation Manager (IM) when selecting Set advanced options . Only when a LockBox has not been created during installation, manually create a LockBox. For more information, see Create a Lockbox on page 139. If you want to use SNMP, ESRS, or LDAP, you must have a LockBox.	Optional	
Configure EMC Secure Remote Support (ESRS)	For more information, see Configuring ESRS connection properties on page 249.	Recommended (to enable Call Home feature)	
Define native users in the system and change the default admin password.	For instructions, see Configure native users on page 141.	Optional	
Define LDAP users (if using ScaleIO with LDAP).	For instructions, see Configure LDAP users on page 142. For detailed information, see <i>EMC ScaleIO User Roles and LDAP Usage Technical Notes</i> .	Optional	
Enable SNMP and configure the SNMP trap receiver	If SNMP was not configured during installation using the advanced settings in the Installation Manager, it can be configured post-deployment.	Optional	

Procedure	Notes	Mandatory?	Done
	For more information, see Configuring SNMP after deployment on page 143.		
Change the Monitor user password		Optional	
Create and map volumes	Creating and mapping volumes can be performed using several management tools. For more information, see Create and map volumes on page 144	Mandatory	

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 ESRS

Enable EMC Secure Remote Support (ESRS) for remote support.

Before you begin

Ensure that a LockBox has already been created and the MDM credentials have been added to it.

Enable the ESRS feature in the `gatewayUser.properties` file.

Note

Detailed information about configuring ESRS is provided in [Before configuring ESRS](#) on page 250.

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_esrs` and edit it as follows:

```
features.enable_esrs=true
```

3. Save and close the file.
4. Add the ESRS Gateway's certificate to the truststore.
 - a. In a web browser, browse to `<ESRS_Gateway_IP_address>:9443`.
 - b. Download the certificate that is displayed, and save it as a file.

c. Log in to REST and get a token. Make a note of the token.

```
curl -k -v --basic --user
admin:<mdm_admin_password> https://
<ScaleIO_Gateway_IP_address>/api/login
```

d. Using REST, add the certificate to truststore.

```
curl -k -v --basic -
uadmin:<token_received_from_previous_command> --form
"file=@<path_to_certificate_file>" https://
<ScaleIO_Gateway_IP_address>/api/trustHostCertificate/Mdm
```

5. Restart the scaleio-gateway service:

- Linux: Run the command `service scaleio-gateway restart`
- Windows: Restart the EMC Gateway service.

6. Register an ESRS gateway:

```
<SioGWTool_PATH> --register_esrs_gateway --scaleio_gateway_ip
<IP_ADDRESS:PORT> --scaleio_gateway_user <USER> --
scaleio_gateway_password <PASSWORD> --esrs_gateway_ip
<IP_ADDRESS> --esrs_gateway_user <USER> --
esrs_gateway_password <PASSWORD> --connect_in_ip <IP_ADDRESS>
```

Configure native users

Configure native users using the CLI.

Use the following SCLI commands to create and modify native ScaleIO users. To run CLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *EMC ScaleIO CLI Reference Guide*. SCLI is installed as part of the MDM component and can be found in the following path:

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

Procedure

1. Log in as an admin user to the CLI:

```
scli --login --username <NAME> --password <PASSWORD>
```

where `<PASSWORD>` is the password used for the first login.

2. (Optional) Change the default admin user password:

```
scli --set_password --old_password <PASSWORD> --new_password
<NEW_PASSWORD>
```

3. Add users:

```
scli --add_user --username <NAME> --user_role <ROLE>
```

A randomly generated password for the created user is returned.

4. (Optional) Log in with the new user and then change its password:

```
scli --login --username <NAME> --password <PASSWORD>
```

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

This is optional because the GUI, CLI, or REST will enforce a password change if user logs in with the original password.

Configure LDAP users

Configure LDAP users if you are using LDAP with ScaleIO.

For more information on using LDAP with ScaleIO, see the *EMC ScaleIO User Roles and LDAP Usage Technical Notes*. To run SCLI commands, you must be logged in. Actual command syntax is operating-system dependent. For more information, see the *EMC ScaleIO CLI Reference Guide*.

Procedure

1. Create Active Directory (AD) groups that correspond to the user roles offered by ScaleIO.
2. Use the CLI to add LDAP service to the MDM:

```
scli --add_ldap_service --ldap_service_uri <URI> --  
ldap_base_dn <LDAP_DN>
```

Note

ScaleIO systems support authentication by up to eight LDAP servers. When multiple LDAP servers are used, add each one separately using this command.

3. Use the CLI to assign an LDAP group to the user role:

```
scli --assign_ldap_groups_to_roles (--ldap_service_id  
<LDAP_SERVICE_ID> | --ldap_service_name <LDAP_SERVICE_NAME>)  
[--administrator_role_dn <LDAP_GROUP_DN>] [--security_role_dn  
<LDAP_GROUP_DN>] [--backend_config_role_dn <LDAP_GROUP_DN>]  
[--frontend_config_role_dn <LDAP_GROUP_DN>] [--  
monitor_role_dn <LDAP_GROUP_DN>]
```

4. Use the CLI to set the system to mixed authentication method, LDAP and native:

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

5. Log in again to apply these changes.

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.

Create and map volumes

Create volumes from devices added to SDS nodes, and then map the volumes to SDC nodes.

You can create and map volumes using various management tools, including the CLI, GUI, vSphere web plug-in, and REST API.

Creating and mapping volumes is necessary 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. For additional details on creating and mapping volumes, see the user documentation.

Creating and mapping volumes using the CLI

Create a volume and map a volume to an SDC using the CLI.

Use the following SCLI commands to create and map volumes. For more details on SCLI command usage, see the *EMC ScaleIO CLI Reference Guide*.

Procedure

1. Create a volume:

```
scli --add_volume --size_gb <SIZE> --volume_name <VOL_NAME> --
protection_domain_name <PD_NAME> --storage_pool_name <SP_NAME>
```

2. Map a volume to an SDC:

```
scli --map_volume_to_sdc --volume_name <VOL_NAME> --sdc_ip
<SDC_IP>
```

Creating and mapping volumes using the GUI

Create a volume and map a volume to an SDC using the GUI.

Before you begin

Ensure that you can authenticate to the GUI with an admin user.

For more information on creating and mapping volumes using the GUI, see the user documentation.

Procedure

1. Open the GUI and log in with credentials:

From the Windows Start menu, click **Start > All Programs > EMC > EMC ScaleIO GUI**.

The ScaleIO GUI is displayed.

2. Open the GUI:

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

- Windows: Click **Start > All Programs > ScaleIO GUI**
3. Create volumes:
 - a. In the **Frontend > Volumes** view, select a Storage Pool to which to add the volume.
 - b. From the **Command** menu or context-sensitive menu, select **Add Volume**.
The **Add Volume** window is displayed.
 - c. To create more than one volume, select **Create multiple volumes** and type the number of volumes in the **Copies** box.
 - d. Type a name in the **Name** box.
 - e. 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.
 - f. Type a number in the **Size** box, representing the volume size in GB (allocation granularity is 8 GB).
 - g. Select **Thick** (default) or **Thin** provisioning options.
 - h. Enable or disable the RAM Read Cache feature by selecting or clearing **Use RAM Read Cache**.
 - i. Click **OK**.
 4. Mapping volumes to an SDC:
 - a. In the **Frontend > Volumes** view, select the volumes.
 - b. 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.
 - c. In the **Select Nodes** panel, select one or more SDCs to which you want to map the volumes.
 - d. 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.

Creating and mapping volumes using the vSphere plug-in

Create a volume and map a volume to an SDC using the vSphere plug-in.

Procedure

1. To open the plug-in, from the vSphere Web Client home tab, click:



2. From the **Storage Pools** screen, click **Actions > Create volumes**.
3. In the **Create Volume** dialog, enter the volume information.

4. To map the volume to ESXs, select **Map volume to ESXs**.
5. In the **Select ESXs** area, select the clusters or ESXs to which this volume should be mapped.

CHAPTER 6

Licensing

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

- [Licensing overview](#) 148
- [Activating entitlements and installing a license file](#) 149
- [License file example](#) 153
- [Error messages](#) 154

Licensing overview

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

Using 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>.

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

You download 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 system, as described in “[Installing the license](#)”.

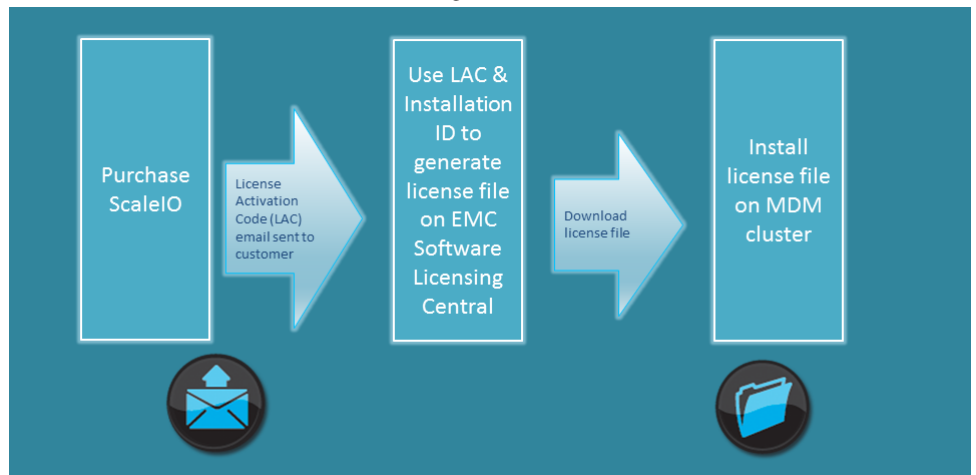
You can view current license information using the CLI or the GUI.

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 11 eLicensing terminology(continued)

Term	Description
EMC Online Support	The EMC online support portal, http://support.emc.com , contains product support

Table 11 eLicensing terminology(continued) (continued)

Term	Description
	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 16 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 ?
PRODUCT LINE: ScaleIO			
EMC SIO SW lic key delivery=CB Product # 436-110-229	0	15	5

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 17 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 12 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.

PART 3

Reference

This part supplies additional information on advanced reference topics. Chapters include:

[Chapter 7, "Manual Deployment"](#)

[Chapter 8, "Advanced Topics"](#)

[Chapter 9, "Maintaining a ScaleIO System"](#)

[Chapter 10, "System Analysis"](#)

[Chapter 11, "Configuring ESRS connection properties"](#)

[Chapter 12, "Common Tasks"](#)

CHAPTER 7

Manual Deployment

This chapter describes how to deploy ScaleIO components in situations when the Installation Manager and VMware deployment wizard are not suitable.

NOTICE

This chapter is for reference purposes only. It is highly recommended to install ScaleIO with the Installation Manager or the VMware deployment wizard. If that is not possible, contact EMC Support.

Topics include:

- [Manual deployment on physical servers](#).....158
- [Manual Deployment of ScaleIO on ESXi Servers](#).....180
- [Manually removing the VIB file from the ESX host](#).....184

Manual deployment on physical servers

This section describes how to deploy ScaleIO manually on physical servers.

It is highly recommended to use the automated Installation Manager as described in [Deploying ScaleIO on Physical Servers](#) on page 65.

Note

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 EMC Support at <https://support.emc.com>.

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

Before deploying SDC or RFCache on Ubuntu, hLinux, Oracle Linux (OL), or CoreOS servers, see the requirements described in [Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers](#) on page 91.

To deploy ScaleIO manually, you must install the following components:

- MDM

This package is for the MDM and Tie Breaker (TB) nodes. You can install MDM in Cluster mode (either 3-node or 5-node) or in Single mode (a single MDM on one node).

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.

- SDS

The SDS is installed on all servers that will contribute storage devices to the ScaleIO system.

- SDC

The SDC is installed on every server that will expose ScaleIO volumes to the application running on it.

- LIA

Install the Lightweight Installation Agent (LIA) on every node in the ScaleIO system. This enables the performance of various maintenance operations (upgrade, uninstall, and others).

- Gateway (optional)

The Gateway can be installed on a Windows or Linux server, within, or external to, the ScaleIO system (though not on an SDC node nor on an SDS on which RFCache is enabled). To use the Gateway, you must also install the LIA on every server.

Note

Before installing the component packages on Linux servers, you must install the GPG key. From the ScaleIO installation folder, run `rpm --import RPM-GPG-KEY-ScaleIO` on every server.

Before installing, ensure that your servers meet the requirements listed in [System requirements](#) on page 18.

Manual installation may require use of SCLI commands. These commands are described, in full, in the *ScaleIO CLI Reference Guide*.

You can deploy as a 3-node or 5-node cluster.

The procedures are described with Linux RHEL/CentOS steps.

The following table describes modifications for deploying on other operating systems:

OS	Modifications required
Ubuntu/hLinux/OL	<p>Before installing ScaleIO, ensure that you have followed the procedures described in Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers on page 91.</p> <ul style="list-style-type: none"> ScaleIO component packages are delivered as TAR files. Before installing, perform the following: <ol style="list-style-type: none"> Untar all the packages: <code>tar -xvf <tar_file></code> This yields SIOB files. Extract the DEB from the SIOB files: <code>siob_extract <siob_file></code> The DEB file is in the following format: <code>EMC-ScaleIO-mdm-2.5-<build>.X.<operating_system>.X.X.x86_64.deb</code> You will use the extracted DEB files for the installation. Some commands are a bit different, noted where applicable.
CoreOS	<p>Before installing ScaleIO, ensure that you have followed the procedures described in Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers on page 91.</p> <ul style="list-style-type: none"> ScaleIO component CoreOS packages are delivered as TAR files. Before installing, perform the following: <ol style="list-style-type: none"> Untar all the packages: <code>tar -xvf <tar_file></code> This yields SIOB files. Extract the BSX from the SIOB files: <code>siob_extract <siob_file></code> The BSX file is in the following format: <code>EMC-ScaleIO-mdm-2.5-<build>.X.CoreOS.x86_64.bsx</code> You will use the extracted BSX files for the installation. Some commands are a bit different, noted where applicable.
Windows	<ul style="list-style-type: none"> Windows package names are in the following format: <code>EMC-ScaleIO-<component>-2.5-<build>.X.msi</code> To restart services, use Services > Restart. Some commands are a bit different, as noted where applicable.

MDM cluster setup

You can set up the MDM cluster in the following modes:

- **3-node cluster**
One Master MDM, one Slave MDM, and one Tie Breaker (TB). This cluster can sustain a single point of failure.
- **5-node cluster**
One Master MDM, two Slave MDMs, and two Tie Breakers. This extended cluster can sustain two points of failure.
- **Single-node cluster**
One Master MDM. Single-node cluster has only one copy of the repository, thus it cannot sustain a failure. It is not recommended to use Single Mode in production systems, except in temporary situations.

An MDM can be installed on any server. For example, one MDM can run on an application server, while the second MDM can run on any other server, for example one of the SDS nodes.

Note

For CLI commands, in cluster mode, use the IP address of the Master MDM or the virtual IP address.

The following is the proper order for creating an MDM cluster:

Order of creating a cluster:

1. Install the MDMs with Manager role (2 in a 3-node, or 3 in a 5-node cluster).
2. Install the MDMs with TB role (1 in a 3-node, or 2 in a 5-node).
3. Create a single-node cluster from one of the Manager MDMs.
4. Add Slave MDMs to the cluster, as Standbys.
5. Add TBs to the cluster, as Standbys.
6. Switch the cluster to 3 or 5-node mode, while designating the standby Slave and standby TB names in the cluster.

Proceed to the section that describes the cluster mode to install:

[“Deploying a 3-node cluster”](#)

[“Deploying a 5-node cluster”](#)

Deploying a 3-node cluster

For this example for a 3-node cluster, we will use the following sample topology:

Table 13 Sample manual installation topology

Role	Management IP	MDM IP	Virtual IP address	Virtual IP NIC	MDM name	Cluster role
MDM 1	10.103.110.152	192.168.1.152	192.168.111.152	eth4	mdm152	Master
MDM 2	10.103.110.153	192.168.1.153	192.168.111.152	eth4	mdm153	Slave

Table 13 Sample manual installation topology (continued)

Role	Management IP	MDM IP	Virtual IP address	Virtual IP NIC	MDM name	Cluster role
TB 1	10.103.110.154	192.168.1.154			tb154	TB
SDS and SDC will be installed on all nodes.						

Installing the packages

Install the ScaleIO packages on servers to be part of a 3-node cluster.

Before you begin

Before deploying ScaleIO on Windows servers, you must run the `Clean_XC_registry.bat` script on every Windows machine to be part of the system. The script is included in the Complete Windows download.

In this section, be sure to use the installation packages and commands that are appropriate for your operating system environment.

Note

If you intend to use the replication feature via RecoverPoint on RHEL 6.x servers, refer to the *ScaleIO Write Splitter for RecoverPoint Technical Notes* before beginning the installation. Replication support is version-specific; see the ESSM for full details.

To install ScaleIO components, perform the following:

Procedure

- Copy the following file to all MDM and TB servers (MDM 1, MDM 2, and TB 1):
 - EMC-ScaleIO-mdm-2.5-<build>.X.elX.x86_64.rpm
- Copy the following files to all servers that will be SDC and/or SDS servers (in our example, all the servers):
 - EMC-ScaleIO-sds-2.5-<build>.X.elX.x86_64.rpm
 - EMC-ScaleIO-sdc-2.5-<build>.X.elX.x86_64.rpm
- Install the MDM and configure the Manager role, by running the following command on MDM 1 and MDM 2:

OS	Command
RHEL/ CentOS /OL	<pre>MDM_ROLE_IS_MANAGER=1 rpm -i <mdm_path.rpm></pre> <p>Example</p> <pre>MDM_ROLE_IS_MANAGER=1 rpm -i EMC-ScaleIO-mdm-2.5-<build>.X.el6.x86_64.rpm</pre>
Ubuntu	<pre>MDM_ROLE_IS_MANAGER=1 dpkg -i <mdm_path.deb></pre>

OS	Command
CoreOS	<pre>MDM_ROLE_IS_MANAGER=1 ./<mdm_path>.bsx</pre>
Windows	<pre>msiexec /i <mdm_path>.msi MDM_ROLE_IS_MANAGER=1</pre> <p>Example</p> <pre>msiexec /i EMC-ScaleIO-mdm-2.5-<build>.X.msi MDM_ROLE_IS_MANAGER=1</pre>

Note

The default MDM credentials are:

- user name: admin
 - password: admin
-

4. Install the MDM on TB 1:

OS	Command
RHEL/ CentOS /OL	<pre>MDM_ROLE_IS_MANAGER=0 rpm -i <mdm_path>.rpm</pre> <p>Example:</p> <pre>MDM_ROLE_IS_MANAGER=0 rpm -i EMC-ScaleIO-mdm-2.5- <build>.X.el6.x86_64.rpm</pre>
Ubuntu	<pre>MDM_ROLE_IS_MANAGER=0 dpkg -i <mdm_path>.deb</pre>
CoreOS	<pre>MDM_ROLE_IS_MANAGER=0 ./<mdm_path>.bsx</pre>
Windows	<pre>msiexec /i <mdm_path>.msi MDM_ROLE_IS_MANAGER=0</pre> <p>Example</p> <pre>msiexec /i EMC-ScaleIO-mdm-2.5-<build>.X.msi MDM_ROLE_IS_MANAGER=0</pre>

5. Install the SDS on every server that will contribute storage drives to the ScaleIO system:

```
rpm -i <sds_path.rpm>
```

Example

```
rpm -i EMC-ScaleIO-sds-2.5-<build>.X.el6.x86_64.rpm
```

Note

Including the IP addresses of as many potential MDMs at this time (up to 8), or of the cluster virtual IP addresses, will make it easier to switch MDMs in the future.

6. Install the SDC on every server that will expose ScaleIO volumes to the application running on it. You can add all the IP addresses of each MDM, or the virtual IP addresses for the MDMs.

Note

Before installing SDC or RFCache on Ubuntu, Oracle Linux (OL), or CoreOS servers, perform the steps described in [Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers](#) on page 91. SDC installation on hLinux is described in the "Installing SDC on a Linux server and connecting it to ScaleIO" section.

Option	Description
RHEL/ CentOS / OL	<pre>MDM_IP=<list of virtual IP or MDM IP addresses> rpm -i <sdc_path.rpm></pre> <p>Example</p> <pre>MDM_IP=192.168.1.152,192.168.1.153,192.168.1.154 rpm -i EMC-ScaleIO-sdc-2.5-<build>.X.el6.x86_64.rpm</pre>
Ubuntu	<pre>MDM_IP=<list of virtual IP or MDM IP addresses> dpkg -i <sdc_path.deb></pre>
CoreOS	<pre>MDM_IP=<list of virtual IP or MDM IP addresses> ./<sdc_path.bsx></pre>
Windows	<pre>msiexec /i <sdc_path.msi> MDM_IP="<list of virtual IP or MDM IP addresses>"</pre>

Option	Description
	<p>Example</p> <pre>msiexec /i EMC-ScaleIO-sdc-2.5-<build>.X.msi MDM_IP="192.168.1.152,192.168.1.153,192.168.1.154"</pre>

All of the components are installed. The next step is to form the MDM management cluster.

Creating the MDM cluster

Procedure

1. Create the MDM single-node cluster on MDM 1, by running the following command on MDM 1:

```
scli --create_mdm_cluster --master_mdm_ip <IP of MDM 1> --
master_mdm_management_ip <Management IP of MDM 1> --
master_mdm_name <Name to give MDM 1> [--cluster_virtual_ip
<IP>] [--master_mdm_virtual_ip_interface <INTF>] --
accept_license
```

Example:

```
scli --create_mdm_cluster --master_mdm_ip 192.168.1.152 --
master_mdm_management_ip 10.103.110.152 --master_mdm_name
mdm152 --cluster_virtual_ip 192.168.111.152 --
master_mdm_virtual_ip_interface eth4 --accept_license
```

Follow these rules:

- Each MDM name must be unique. Ensure that each server on which MDM is installed has a unique hostname.
 - Each virtual IP address must be related to one interface.
 - When listing multiple virtual IP/interface combinations, the list of interfaces must be in the same order as the list of virtual IP addresses.
2. If prompted, press Y to continue.
 3. Log in and change the MDM password, by running the following commands on MDM 1:

```
scli --login --username admin --password admin
```

```
scli --set_password --old_password admin --new_password
Scaleio123
```

The login password is changed from admin to Scaleio123.

4. Log in with the new credentials:

```
scli --login --username admin --password Scaleio123
```

5. Add MDM 2 to the MDM cluster as a Standby MDM, by running the following command on MDM 1:

```
scli --add_standby_mdm --new_mdm_ip <IP of MDM 2> --mdm_role
manager --new_mdm_management_ip <Management IP of MDM 2> --
new_mdm_virtual_ip_interface <INTF> --new_mdm_name <Name to
give MDM 2>
```

Example:

```
scli --add_standby_mdm --new_mdm_ip 192.168.1.153 --mdm_role
manager --new_mdm_management_ip 10.103.110.153 --
new_mdm_virtual_ip_interface eth4 --new_mdm_name mdm153
```

6. Add TB 1 to the MDM cluster as a Tie Breaker, by running the following command on MDM 1:

```
scli --add_standby_mdm --new_mdm_ip <IP of TB 1> --mdm_role
tb --new_mdm_name <Name to give TB 1>
```

Example:

```
scli --add_standby_mdm --new_mdm_ip 192.168.1.154 --mdm_role
tb --new_mdm_name tb154
```

7. (Optional) Query the cluster to verify that it's running in single mode, by running the following command on MDM 1:

```
scli --query_cluster
```

Output, similar to the following, is displayed:

```
Cluster: Mode: 1_node
```

8. Switch to 3-node cluster mode, by running the following command on MDM 1:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_name <MDM 2 name> --add_tb_name <TB 1 name>
```

Example:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_name mdm153 --add_tb_name tb154
```

9. (Optional) Query the cluster to verify that it's running in cluster mode, by running the following command on MDM 1:

```
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
  Virtual IPs: 192.168.111.152
Master MDM:
  Name: MDM152, ID: 0x7d76db0d2ad840b0
  IPs: 192.168.1.152, Management IPs: 10.103.110.152,
  Port: 9011, Virtual IP interfaces: eth4
  Version: 2.0.13000
Slave MDMs:
  Name: MDM153, ID: 0x41788ea96626fe21
  IPs: 192.168.1.153, Management IPs: 10.103.110.153,
  Port: 9011, Virtual IP interfaces: eth4
  Status: Normal, Version: 2.0.13000
Tie-Breakers:
  Name: MDM154, ID: 0x3d8263b64fcf2642
  IPs: 192.168.1.154, Port: 9011
  Status: Normal, Version: 2.0.13000
```

After completion, see [“Post-installation considerations”](#).

Deploying a 5-node cluster

For this example for a 5-node cluster, we will use the following sample topology:

Table 14 Sample manual installation topology

Role	Management IP	MDM IP	Virtual IP address	Virtual IP NIC	MDM name	Cluster role
MDM 1	10.103.110.152	192.168.1.152	192.168.111.152	eth4	mdm152	Master
MDM 2	10.103.110.153	192.168.1.153	192.168.111.152	eth4	mdm153	Slave 1
MDM 3	10.103.110.154	192.168.1.154	192.168.111.152	eth4	mdm154	Slave 2
TB 1	10.103.110.155	192.168.1.155			tb155	TB 1
TB 2	10.103.110.156	192.168.1.156			tb156	TB 2
SDS and SDC will be installed on all nodes.						

Installing the packages

Install the ScaleIO packages on servers to be part of a 5-node cluster.

Before you begin

Before deploying ScaleIO on Windows servers, you must run the `Clean_XC_registry.bat` script on every Windows machine to be part of the system. The script is included in the Complete Windows download.

In this section, be sure to use the installation packages and commands that are appropriate for your operating system environment.

Note

If you intend to use the replication feature via RecoverPoint on RHEL 6.x servers, refer to the *ScaleIO Write Splitter for RecoverPoint Technical Notes* before beginning the installation. Replication support is version-specific; see the ESSM for full details.

To install ScaleIO components, perform the following:

Procedure

1. Copy the following file to all MDM and Tie Breaker (TB) servers (MDM 1, 2, 3 and TB 1 and 2):
 - a. EMC-ScaleIO-mdm-2.5-<build>.X.elX.x86_64.rpm
2. Copy the following files to all servers that will be SDC and/or SDS servers (in our example, all the servers):
 - EMC-ScaleIO-sds-2.5-<build>.X.elX.x86_64.rpm
 - EMC-ScaleIO-sdc-2.5-<build>.X.elX.x86_64.rpm
3. Install the MDM and configure the Manager role, by running the following commands on the Master (MDM 1) and Slave MDMs (MDM 2 and MDM 3):

OS	Command
RHEL/ CentOS /OL	<pre>MDM_ROLE_IS_MANAGER=1 rpm -i <mdm_path.rpm></pre> <p>Example</p> <pre>MDM_ROLE_IS_MANAGER=1 rpm -i EMC-ScaleIO-mdm-2.5-<build>.X.el6.x86_64.rpm</pre>
Ubuntu	<pre>MDM_ROLE_IS_MANAGER=1 dpkg -i <mdm_path.deb></pre>
CoreOS	<pre>MDM_ROLE_IS_MANAGER=1 /.<mdm_path.bsx></pre>
Windows	<pre>msiexec /i <mdm_path.msi> MDM_ROLE_IS_MANAGER=1</pre> <p>Example</p> <pre>msiexec /i EMC-ScaleIO-mdm-2.5-<build>.X.msi MDM_ROLE_IS_MANAGER=1</pre>

Note

The default MDM credentials are:

- user name: admin
- password: admin

4. Install the MDM on the Tie Breaker MDMs (TB 1 and TB 2):

OS	Command
RHEL/ CentOS /OL	<pre>MDM_ROLE_IS_MANAGER=0 rpm -i <mdm_path.rpm></pre> <p>Example:</p> <pre>MDM_ROLE_IS_MANAGER=0 rpm -i EMC-ScaleIO-mdm-2.5- <build>.X.el6.x86_64.rpm</pre>
Ubuntu	<pre>MDM_ROLE_IS_MANAGER=0 dpkg -i <mdm_path.deb></pre>
CoreOS	<pre>MDM_ROLE_IS_MANAGER=0 ./<mdm_path.bsx></pre>
Windows	<pre>msiexec /i <mdm_path.msi> MDM_ROLE_IS_MANAGER=0</pre> <p>Example</p> <pre>msiexec /i EMC-ScaleIO-mdm-2.5-<build>.X.msi MDM_ROLE_IS_MANAGER=0</pre>

5. Install the SDS on every server that will contribute storage drives to the ScaleIO system:

```
rpm -i <sds_path.rpm>
```

Example

```
rpm -i EMC-ScaleIO-sds-2.5-<build>.X.el6.x86_64.rpm
```

Note

Including the IP addresses of as many potential MDMs at this time (up to 8), or of the cluster virtual IP addresses, will make it easier to switch MDMs in the future.

6. Install the SDC on every server that will expose ScaleIO volumes to the application running on it. You can add all the IP addresses of each MDM, or the virtual IP addresses for the MDMs.

Note

Before installing SDC or RFcache on Ubuntu, Oracle Linux (OL), or CoreOS servers, perform the steps described in [Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers](#) on page 91. SDC installation on hLinux is described in the "Installing SDC on a Linux server and connecting it to ScaleIO" section.

Option	Description
RHEL/ CentOS / OL	<pre>MDM_IP=<list of virtual IP or MDM IP addresses> rpm -i <sd_c_path.rpm></pre> <p>Example</p> <pre>MDM_IP=192.168.1.152,192.168.1.153,192.168.1.154 rpm -i EMC-ScaleIO-sdc-2.5-<build>.X.el6.x86_64.rpm</pre>
Ubuntu	<pre>MDM_IP=<list of virtual IP or MDM IP addresses> dpkg -i <sd_c_path.deb></pre>
CoreOS	<pre>MDM_IP=<list of virtual IP or MDM IP addresses> ./<sd_c_path.bsx></pre>
Windows	<pre>msiexec /i <sd_c_path.msi> MDM_IP="<list of virtual IP or MDM IP addresses>"</pre> <p>Example</p> <pre>msiexec /i EMC-ScaleIO-sdc-2.5-<build>.X.msi MDM_IP="192.168.1.152,192.168.1.153,192.168.1.154"</pre>

All of the components are installed. The next step is to form the MDM management cluster.

Creating the MDM cluster

Procedure

1. Create the MDM single-node cluster on MDM 1, by running the following command on MDM 1:

```
scli --create_mdm_cluster --master_mdm_ip <IP of MDM 1> --  
master_mdm_management_ip <Management IP of MDM 1> --  
master_mdm_name <Name to give MDM 1> [--cluster_virtual_ip
```

```
<IP>] [--master_mdm_virtual_ip_interface <INTF>] --
accept_license
```

Example:

```
scli --create_mdm_cluster --master_mdm_ip 192.168.1.152 --
master_mdm_management_ip 10.103.110.152 --master_mdm_name
mdm152 --cluster_virtual_ip 192.168.111.152 --
master_mdm_virtual_ip_interface eth4 --accept_license
```

Note

Each MDM name must be unique. Ensure that each server on which MDM is installed has a unique hostname.

2. Log in and change the MDM password, by running the following commands on MDM 1:

```
scli --login --username admin --password admin
```

```
scli --set_password --old_password admin --new_password
Scaleio123
```

The login password is changed from admin to Scaleio123.

3. Log in with the new credentials:

```
scli --login --username admin --password Scaleio123
```

4. Add the Slave MDMs (MDM 2 and MDM 3) to the MDM cluster as Standby MDMs, by running the following commands on MDM 1:

```
scli --add_standby_mdm --new_mdm_ip <IP of MDM X> --mdm_role
manager --new_mdm_management_ip <Management IP of MDM X> --
new_mdm_virtual_ip_interface <INTF> --new_mdm_name <Name to
give MDM X>
```

Examples:

```
scli --add_standby_mdm --new_mdm_ip 192.168.1.153 --mdm_role
manager --new_mdm_management_ip 10.103.110.153 --
new_mdm_virtual_ip_interface eth4 --new_mdm_name mdm153
```

```
scli --add_standby_mdm --new_mdm_ip 192.168.1.154 --mdm_role
manager --new_mdm_management_ip 10.103.110.154 --new_mdm_name
mdm154
```

5. Add TB 1 and TB 2 to the MDM cluster as a Tie Breaker, by running the following commands on MDM 1:

```
scli --add_standby_mdm --new_mdm_ip <IP of TB X> --mdm_role
tb --new_mdm_name <Name to give TB X>
```

Examples:

```
scli --add_standby_mdm --new_mdm_ip 192.168.1.155 --mdm_role
tb --new_mdm_name tb155
```

```
scli --add_standby_mdm --new_mdm_ip 192.168.1.156 --mdm_role
tb --new_mdm_name tb156
```

6. (Optional) Query the cluster to verify that it's running in single mode, by running the following command on MDM 1:

```
scli --query_cluster
```

Output, similar to the following, is displayed:

```
Cluster: Mode: 1_node
```

7. Switch to 5-node cluster mode, by running the following command on MDM 1:

```
scli --switch_cluster_mode --cluster_mode 5_node --
add_slave_mdm_name <MDM 2 name>,<MDM 3 name> --add_tb_name
<TB 1 name>,<TB 2 name>
```

Example:

```
scli --switch_cluster_mode --cluster_mode 5_node --
add_slave_mdm_name mdm153,mdm154 --add_tb_name tb155,tb156
```

8. (Optional) Query the cluster to verify that it's running in cluster mode, by running the following command on MDM 1:

```
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
Virtual IPs: 192.168.111.152
Master MDM:
Name: mdm6, ID: 0x1cb79cb46ed13200
IPs: 192.168.247.6, 172.16.247.6, Management IPs: 10.136.247.6,
Port: 9011, Virtual IP interfaces: eth4
Version: 2.0.13000
Slave MDMs:
Name: mdm8, ID: 0x67d759a83ec90d42
```

```

IPs: 172.16.247.8, 192.168.247.8, Management IPs: 10.136.247.8,
Port: 9011, Virtual IP interfaces: eth4
Status: Normal, Version: 2.0.13000
Name: mdm7, ID: 0x4d9cb20c720f7481
IPs: 192.168.247.7, 172.16.247.7, Management IPs: 10.136.247.7,
Port: 9011, Virtual IP interfaces: eth4
Status: Normal, Version: 2.0.13000
Tie-Breakers:
Name: mdm9, ID: 0x64bd85740ae9b2a4
IPs: 172.16.247.9, 192.168.247.9, Port: 9011
Status: Normal, Version: 2.0.13000
Name: mdm10, ID: 0x593656bc4ef59083
IPs: 172.16.247.10, 192.168.247.10, Port: 9011
Status: Normal, Version: 2.0.13000

```

Post-installation considerations

This section describes various post-installation considerations. Proceed to the section that matches your operating system environment.

Linux

Each installed component creates subdirectories for its use in the following directory:

/opt/emc/scaleio/<component>/

For example: /opt/emc/scaleio/mdm/

Note

If you are running Veritas Cluster Server (VCS) software, see [ScaleIO support on Linux servers running Veritas Cluster Server \(VCS\) software](#) on page 83 for post-installation instructions.

Installing RCache - Linux

To enable RCache, where the RCache devices provide acceleration, you need to install it, then enable it.

Every SDS that provides acceleration must have at least one RCache device on it. For a SP with RCache enabled, you must select at least one RCache device for each SDS whose devices are part of that SP.

Before installing RCache on Ubuntu, Oracle Linux, or CoreOS servers, see the requirements described in [Deploying on CoreOS, hLinux, Oracle Linux, or Ubuntu servers](#) on page 91.

Procedure

1. Install RCache on every SDS server that is to provide acceleration.

RCache is enabled with the xcache package:

```
rpm -i EMC-ScaleIO-xcache-2.5-<build>.X.elX.x86_64.rpm
```

2. Add the RCache device, by running the following command on the Master MDM:

```
scli --add_sds_rcache_device --sds_name <SDS_NAME> --
rcache_device_path <PATH> --rcache_device_name <DEVICE_NAME>
```

Example:

```
scli --add_sds_rfcache_device --sds_name sds_192 --
rfcache_device_path /dev/sdh --rfcache_device_name sdh
```

3. Enable RFCache on the SP, by running the following command on the Master MDM:

```
scli --set_rfcache_usage --use_rfcache --
protection_domain_name <PD_NAME> --storage_pool_name <SP_NAME>
```

Example:

```
scli scli --set_rfcache_usage --use_rfcache --
protection_domain_name pd_1 --storage_pool_name sp_1
```

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."

Installing the Gateway - Linux

The ScaleIO Gateway includes the REST Gateway, the Installation Manager (IM), and the SNMP trap sender and ESRs functionality.

To install the ScaleIO Gateway, follow the instructions described in ["Preparing the Installation Manager and the Gateway"](#).

CoreOS servers do not support Gateway installation. Deploy a Gateway on another Windows or Linux server.

The Gateway relies on the presence of the LIA on every server. To install the LIA, see the following section.

Installing the LIA - Linux

You can install the LIA on Linux servers.

Run the procedures here on every server where the LIA is being installed.

Procedure

1. Copy the LIA package to the server:

```
EMC-ScaleIO-lia-2.5-<build>.X.elX.x86_64.rpm
```

2. Install the LIA component:

```
TOKEN=<password> rpm -i <Full RPM file path>
```

The password must meet the following criteria:

- a. Between 6 and 31, ASCII-printable characters
- b. No blank spaces
- c. Include at least 3 of the following groups: [a-z], [A-Z], [0-9], special chars (!@#\$...)

For example:

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

3. 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 hexadecimal number that is displayed into the file created in the previous step.

- d. Restart the LIA service:

```
pkill lia
```

The LIA component requires security configuration, as described in the *EMC ScaleIO Security Configuration Guide*.

After installing, you need to provision storage. See [“Provisioning storage example”](#).

After you finish

It is highly recommended to run the ScaleIO system analysis to analyze the ScaleIO system immediately after deployment, before provisioning volumes, and before using the system in production. For more information, see [System analysis overview](#) on page 240.

After installing, you need to provision storage. See [“Provisioning storage example”](#).

Note

For optimal performance in environments with more than 60,000 IOPS, see the *ScaleIO Performance Fine-Tuning Technical Notes*.

Windows

Each installed component creates subdirectories for its use in the following directory:

`c:\Program files\emc\scaleio\<component>\`

For example: `c:\Program files\emc\scaleio\mdm\`

Installing RCache - Windows

To enable RCache, where the RCache devices provide acceleration, you need to install it (the xcache package), then enable it.

When installing xcache, a server restart is necessary.

Every SDS that provides acceleration must have at least one RCache device on it. For a SP with RCache enabled, you must select at least one RCache device for each SDS whose devices are part of that SP.

Procedure

1. Install RCache on every SDS server that is to provide acceleration.

RCache is enabled with the xcache package:

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

Note

After installation completes, you must restart the server.

2. Add the RCache device, by running the following command on the Master MDM:

```
scli --add_sds_rfcache_device --sds_name <SDS_NAME> --  
rfcache_device_path <PATH>
```

Example

```
scli --add_sds_rfcache_device --sds_name sds_192 --  
rfcache_device_path j
```

Note

The RFlcache device path is represented by a drive letter.

3. Enable RFlcache on the SP, by running the following command on the Master MDM:

```
scli --set_rfcache_usage --use_rfcache --  
protection_domain_name <PD_NAME> --storage_pool_name <SP_NAME>
```

Example

```
scli scli --set_rfcache_usage --use_rfcache --  
protection_domain_name pd_1 --storage_pool_name sp_1
```

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."

Installing the Gateway - Windows

The ScaleIO Gateway includes the REST Gateway, the Installation Manager (IM), and the SNMP trap sender functionality.

To install the ScaleIO Gateway, follow the instructions described in ["Preparing the Installation Manager and the Gateway"](#).

Installing the LIA - Windows

You can install the LIA on Windows servers.

Run the procedures here on every server where the LIA is being installed.

Procedure

1. Install the LIA component:

```
msiexec /i EMC-ScaleIO-lia-2.5-<build>.X.msi TOKEN=<password>
```

The password must meet the following criteria:

- a. Between 6 and 31, ASCII-printable characters
 - b. No blank spaces
 - c. 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:

```
C:\Program Files\EMC\scaleio\lia\cfg\installation_id.txt
```

- b. Log in, then query the MDM for the installation ID by running the following command:

```
scli --query_all
```

As part of the output, the installation ID is displayed.

- c. Copy the installation ID into the new file, and add a blank line under it.
- d. Restart the LIA service:

```
EMC scaleio LIA service
```

The LIA component requires security configuration, as described in the *ScaleIO Performance Fine-Tuning Technical Notes*.

Note

For optimal performance in environments with more than 60,000 IOPS, see the *Fine-Tuning ScaleIO Performance Technical Notes*.

Provisioning storage example

This section shows a simple illustration of how to use the CLI to add capacity and create and map volumes. You can also use the GUI, REST API, or the VMware plug-in. The IP addresses in the example are taken from the 3-node cluster manual installation example, described in [“Deploying a 3-node cluster”](#), but the concepts apply to 5-node clusters, too.

Adding capacity

This section describes how to add capacity. For our example, we will use the following names:

Table 15 Provisioning storage example

Item to add	Name
Protection Domain	pd1
Storage Pool	sp1
Devices	/dev/sdb, /dev/sdc, and /dev/sdd

All scli commands are run from the Master MDM (MDM 1), unless otherwise noted.

Procedure

1. Define a Protection Domain, `pd1`, by running the following command, from the Master MDM (MDM 1):

```
scli --add_protection_domain --protection_domain_name pd1
```

Note

To run scli commands from another server, add `--mdm_ip <IP MDM 1>, <IP MDM 2>` to each command line.

2. Add a Storage Pool, `sp1`, by running the following command:

```
scli --add_storage_pool --protection_domain_name pd1 --
storage_pool_name sp1
```

3. Add an SDS on MDM 1, and its devices to `pd1`, and name it `sds152`, by running the following command:

```
scli --add_sds --sds_ip <sds IP>
--device_path <sds devices>
--protection_domain_name <protection domain name>
--storage_pool_name <storage pool name>
--sds_name <sds name>
```

Example:

```
scli --add_sds --sds_ip 192.168.1.152 --device_path /dev/
sdb,/dev/sdc,/dev/sdd --storage_pool_name sp1 --
protection_domain_name pd1 --sds_name sds152
```

Note

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.

To add devices on a Windows-based SDS, you can use either a drive letter or a file for the SDS. For more information, see [“Adding devices to SDS nodes on Windows servers”](#).

4. Add an SDS on MDM 2, and its devices to `pd1`, and name it `sds153`, by running the following command:

```
scli --add_sds --sds_ip 192.168.1.153 --device_path /dev/
sdb,/dev/sdc,/dev/sdd --storage_pool_name sp1 --
protection_domain_name pd1 --sds_name sds153
```

5. Add an SDS on TB 1, and its devices to `pd1`, add meaningful names to the devices (optional), and name the SDS `sds154`, by running the following command:

```
scli --add_sds --sds_ip 192.168.1.154 --device_path /dev/
sdb,/dev/sdc,/dev/sdd --device_name
node_12_dev1,node_12_dev5,node_12_dev7 --storage_pool_name
sp1 --protection_domain_name pd1 --sds_name sds154
```

Creating and mapping volumes

This section describes how to create and map volumes.

Procedure

1. Define a 2500 GB volume, and name it `vol1`, by running the following command:

```
scli --add_volume --volume_name <volume name>
--protection_domain_name <protection domain name>
--storage_pool_name <storage pool name> --size_gb <size>
```

Example:

```
scli --add_volume --volume_name vol1 --protection_domain_name
pd1 --storage_pool_name sp1 --size_gb 2500
```

2. Map `vol1` to the SDC on MDM 1, by running the following command:

```
scli --map_volume_to_sdc --volume_name vol1 --sdc_ip <IP MDM
1>
```

Example:

```
scli --map_volume_to_sdc --volume_name vol1 --sdc_ip
192.168.1.152
```

Configuring spare capacity

This section describes how to configure spare capacity for your system. This step is highly recommended in any system with less than 10 SDSs or Fault Sets.

Modify the spare capacity, by running the following command on MDM 1:

```
scli --modify_spare_policy --protection_domain_name pd1 --
storage_pool_name spl --spare_percentage 34 [--i_am_sure]
```

The number 34 represents the percentage of total capacity set aside to ensure data integrity during server failures. The percentage is derived by $1/(\text{number of SDS})$, which yields the recommended percentage for less than 10 balanced servers. For more information, see “Modifying spare policy” in the CLI Reference Guide.

Manual Deployment of ScaleIO on ESXi Servers

You can deploy ScaleIO on ESXi servers manually.

It is highly recommended to deploy using the vSphere plug-in deployment wizard. In addition, familiarize yourself with the ESX vStorage APIs for Array Integration (VAAI) features supported by the system, as described in the Architecture chapter of your system's Deployment Guide and User Guide.

This section describes tasks that are necessary in cases where this is not possible or desirable, such as for scripting.

After manual deployment on ESX, you must perform manual memory allocation on the SVM, as described in [“SVM manual memory allocation”](#).

Before starting, obtain the IP addresses for all of the nodes to be installed, and an additional IP address for the MDM cluster, as described in the following table:

Table 16 IP address table for manual deployment on ESXi servers

Number of networks	IP addresses required
1 network (management network)	<ul style="list-style-type: none"> 1 for vmnic 1 for SVM
2 networks	<ul style="list-style-type: none"> 2 for vmnic (1 for mgmt, 1 for data) 2 for SVM (1 for mgmt, 1 for data)
3 networks	<ul style="list-style-type: none"> 3 for vmnic (1 for mgmt, 2 for data) 3 for SVM (1 for mgmt, 2 for data)

Deploying the ScaleIO virtual machine (SVM)

This topic describes how to deploy the SVM and prepare it for installation of ScaleIO. This is necessary when the SDC installed on ESX, but all other components are installed on ScaleIO virtual machines (SVM).

Note

Before deploying, ensure that networks for each ESX to be part of the system have a properly defined Port Group and VMkernel.

To deploy the SVM on the ESX, perform the following:

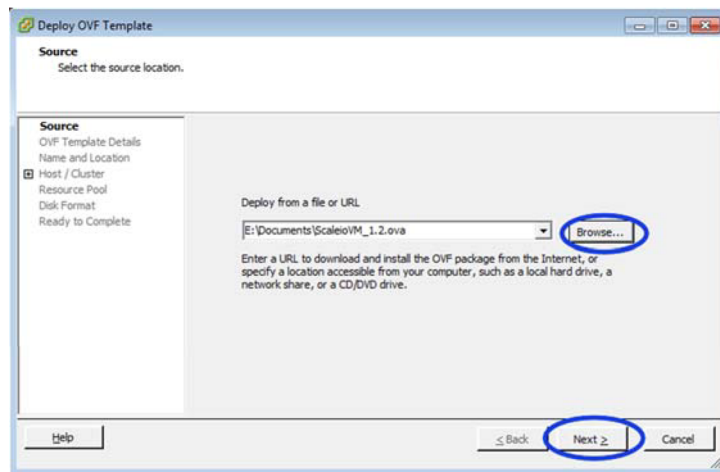
Procedure

1. Download the installation software from the EMC Support site (<https://support.emc.com>).

If necessary, extract the files to an installation folder.

2. Download the following document: <http://www.vmware.com/files/pdf/techpaper/vmware-multipathing-configuration-software-iSCSI-port-binding.pdf>.
3. Use the vSphere client to connect to a vCenter.
4. Select **File > Deploy OVF Template**.

The Deploy OVF Template wizard appears.



5. Enter the full path to the OVA, and click **Next**.
6. Assign a unique name to the VM, accept all remaining default values, and click **Finish**.
7. Clone the SVM to each ESX host, making sure to assign a unique name to each VM.
8. On each ESX physical host, configure the network and NTP (and DNS, if necessary).

- a. Using the console, start the SVM, and login.

The default user name is `root`, and the default password is `admin`.

- b. Configure the network:

- IP address
- Gateway address
- DNS
- NTP server

Note

When using the OVA, the installation packages can be found in `/root/install/`.

9. On each SVM in the system, add the local disk devices as RDM devices. Perform this task by changing the SVM settings.

10. Configure each ESX host to enable *Start and Stop Virtual Machines with the system*, and set the SVM to do *Automatic Startup*.
11. Install the GPG key. From the ScaleIO installation folder, run `rpm --import RPM-GPG-KEY-ScaleIO` on every server.
12. On each SVM in the system, install the relevant ScaleIO components, as described in [“Manual deployment on physical servers”](#).

The packages are located under `/root/install`

Note

After manual deployment on ESX, you must perform manual memory allocation on the SVM, as described in [“SVM manual memory allocation”](#).

Configuring the UUID on virtual machines

To configure the virtual machine, you must enable the UUID attribute on every virtual machine, by performing the following:

Procedure

1. Start the vSphere Client, and log in to a vCenter Server.
2. Select **Virtual Machines and Templates**.
3. From the **Virtual Machines** tab, right-click the virtual machine, and choose **Power > Power Off**.

The virtual machine powers off.

4. Right-click the virtual machine, and click **Edit Settings**.
5. From the **Options** tab, select **Advanced > General** in the settings column.
6. Click **Configuration Parameters**.

The **Configuration Parameters** window appears.

7. Click **Add Row**, and enter the following:
 - a. In the **Name** column, enter `disk.EnableUUID`.
 - b. In the **Value** column, enter `TRUE`.
8. Click **OK**.
9. Power on the virtual machine.
10. Repeat this procedure for every SVM in the ScaleIO system.

Installing the SDC directly on an ESX host

To install the SDC component directly on an ESXi host, perform the following:

Procedure

1. Set the acceptance level of your host to `PartnerSupported`, by entering:

```
esxcli software acceptance set --level=PartnerSupported
```

2. Install the SDC VIB, by entering the following:

```
esxcli software vib install -d <full_path_to_VIB>
```

where `<full_path_to_VIB>` is the full path to the `sdc-2.5-<build>.X-esx5.5.zip` file.

Note

To install without changing the acceptance level, use this command instead:
`esxcli software vib install -f -d <VIB location>`

3. Reboot the ESX host.

The SDC will not automatically boot up at this point.

Update the SDC GUID and MDM IP address parameters:

- a. GUID, for example: 12345678-90AB-CDEF-1234-567890ABCDEF

Note

Including the IP addresses of as many potential manager MDMs as possible at this time will make it easier to switch MDMs in the future. You can add a total of eight IP addresses.

- b. MDM IP addresses.

You can define multiple MDM clusters, each with multiple IP addresses.

- Separate the IP addresses of the same MDM cluster with a “,” symbol.
- Separate multiple MDM clusters with the “+” symbol.

In the following example, there are two MDM clusters. The first has two IP addresses and the second has only one:

```
10.20.30.40,50.60.70.80+11.22.33.44
```

4. To update the SDC parameters, enter the following command:

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=<GUID> IoctlMdmIPStr=<MDM_IPs>"
```

where:

- a. `GUID` is the GUID (maximum of 1024 digits)
- b. `MDM_IPs` is the IP addresses, as described earlier

5. Back up the parameters, by entering the following:

```
/bin/auto-backup.sh
```

6. Load the SDC module, by entering the following:

```
esxcli system module load -m scini
```

Optimizing the guest operating system performance

The configuration of the guest operating system affects the performance of the system. The ScaleIO VM (SVM) is optimized. If the storage VM that is hosting the

SDS and MDM roles is not supplied by ScaleIO, refer to the following guideline and procedures.

Modify the device scheduler (Linux guests only)

The following command modifies the I/O scheduler of the devices. It should be run for each device that exposes a ScaleIO volume inside the guest virtual machine.

Command

```
echo noop > /sys/block/<device_name>/queue/scheduler
```

Example

```
echo noop > /sys/block/sds/queue/scheduler
```

Modifying the SCSI controller type of the guest virtual machine

To configure the SCSI controller type to use a Paravirtual SCSI (PVSCSI) adapter, perform the following steps for each guest using ScaleIO storage.

Note

Before performing the following steps, consult VMware documentation to verify that your guest operating system supports Paravirtual SCSI adapters.

Procedure

1. Power off the guest virtual machine.
2. Open the vSphere client, and right click on the guest.
3. Select **Edit Settings**.
4. Select **SCSI Controller 0**.
5. Click **Change Type**.
6. Select VMware **Paravirtual**.
7. Click **OK**.

Manually removing the VIB file from the ESX host

Procedure

1. In an SSH session, check the version of VMware ESXi installed on the host by typing this command:

```
esxcli system version get
```

Example of output:

```
Product: VMware ESXi
Version: 5.5.0
Build: Releasebuild-2068190
Update: 2
```

In the above example, the version is 5.5.0.

2. Type the command that corresponds to the version installed on the host. For example:

- a. For ESXi 5.5:

```
esxcli software vib remove --vibname scaleio-sdc-esx5.5
```

- b. For ESXi 6.0:

```
esxcli software vib remove --vibname scaleio-sdc-esx6.0
```

Example of output:

```
Removal Result
Message: The update completed successfully, but the system
needs to be rebooted for the changes to be effective.
Reboot Required: true
VIBs Installed:
VIBs Removed: VMware_bootbank_scaleio-sdc-esx5.5_1.32-201.0
VIBs Skipped:
```

3. Open the `esx.conf` file for editing by typing this command:

```
vi /etc/vmware/esx.conf
```

- a. Search for the line containing the `scini` text (use key sequence `/scini`).
- b. Delete the line using either the `<X>` key or `<D>` to delete to full line.
- c. Save and exit using the key sequence: `<ESC>+<:wq!>` and `<Enter>`.

CHAPTER 8

Advanced Topics

This chapter describes advanced ScaleIO installation and upgrade topics. Topics include:

- [Advanced Gateway topics](#)..... 188
- [Advanced Installation Manager topics](#).....193
- [Advanced vSphere plug-in topics](#).....213
- [Deploy ScaleIO in a 2-layer environment](#).....217
- [Prepare the servers for deployment](#)..... 219

Advanced Gateway topics

This section includes advanced procedures for Gateway installation.

Enabling and disabling Gateway components

You can use the Features Enabler to enable and disable each Gateway feature. The Features Enabler is activated by setting the values in the Features Enabler section of the `gatewayUser.properties` file.

- Enable SNMP (default: false)

To enable, set `features.enable_snmp= true`

- Enable Gateway (default: true)

To disable, set `features.enable_gateway=false`

You can disable the use of the default port, 443, by setting both this property and the `features.enable_IM` to false.

- Enable Installation Manager (default: true)

To disable, set `features.enable_IM=false`

You can disable the use of the default port, 443, by setting both this property and the `features.enable_gateway` to false.

- Enable ESRS (default: false)

To enable, set `features.enable_esrs=false true`

- Disable local user login (default: false)

To restrict login to LDAP users only, set

```
gateway-admin.disable.local.login=true
```

- Enable adding SDS devices that have a signature from a previous ScaleIO system (default: false). This can happen in the following cases:

- Add SDS devices that were used in a previous ScaleIO system.
- Extending a system with an SDS whose devices were used in a previous system.
- Extending a system by adding SDSs to a node that already has an SDS installed.

To enable this feature, set `add.sds.with.force.on.next.run=true`

This setting is valid for one deployment only. After being marked completed, it reverts to false.

To configure the properties, perform the following:

Procedure

1. Use a text editor to open the `gatewayUser.properties` file, located in the following directory on the Installation Manager/Gateway server:

Gateway installed on	Location of gatewayUser.properties file
Windows	C:\Program Files\EMC\ScaleIO\Gateway\webapps\ROOT\WEB-INF\classes\
Linux	/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes

2. Edit the file with the desired changes.
 3. Save and close the file.
 4. Restart the scaleio-gateway service:
 - a. Windows:
Restart the EMC ScaleIO Gateway service.
 - b. Linux
Type the command `service scaleio-gateway restart`.
- Configuration is complete.

Using a custom Java configuration for the Gateway

To use a custom Java configuration, use the following command to install the Gateway, in place of the normal command:

```
GATEWAY_ADMIN_PASSWORD=<new_GW_admin_password>
SIO_GW_JAVA=<path_to_java> rpm -U /tmp/EMC-ScaleIO-gateway-2.5-
<build>.X.x86_64.rpm
```

where *<new_GW_admin_password>* is a password that you define to access the Installation Manager administration commands and web interface, and *<path_to_java>* is the path to the desired Java files.

Installing the Gateway without assigning an admin password

You can install the ScaleIO Gateway without adding the `GATEWAY_ADMIN_PASSWORD` variable.

To add this variable later, perform one of the following:

Procedure

1. Use a REST URI command, as described in the *ScaleIO User Guide*.
2. Edit the user properties file, and then restart the scaleio-gateway service, as described in the *ScaleIO User Guide*.

Certificate management for ScaleIO Gateway

This section explains how to replace the ScaleIO gateway's self-signed security certificate with your organization's "trusted" certificate, and how to create a new "trusted" certificate. The ScaleIO gateway automatically creates its own self-signed security certificate when it is installed or upgraded. If your organization has no special security certificate requirements, you can keep working with the default certificate.

Replacing the default self-signed security certificate with your own trusted certificate

You can create your own trusted certificate, and then replace the default certificate with the one that you created.

Procedure

1. Find the location of `keytool` on your server, and open it.

It is a part of the Java (JRE or JDK) installation on your server, in the `bin` directory. For example:

- a. `C:\Program Files\Java\jdk1.7.0_25\bin\keytool.exe`
- b. `/usr/bin/keytool`

2. Generate your RSA private key, by typing this command:

```
keytool -genkey -alias <YOUR_ALIAS> -keyalg RSA -keystore
<PATH_TO_NEW_KEYSTORE_FILE>
```

- a. If you want to define a password, add the following parameters to the command. Use the same password for both parameters.

```
-storepass <KEYSTORE_PASSWORD> -keypass <KEYSTORE_PASSWORD>
```

Note

Specify a directory outside the ScaleIO Gateway installation directory for the newly created keystore file. This will prevent it from being overwritten when the ScaleIO Gateway is upgraded or reinstalled.

3. If you already have a Certificate Signing Request (CSR), skip this step.

If you need a CSR, generate one by typing the following command. (If you did not define a keystore password in the previous step, omit the password flags.)

```
keytool -certreq -keyalg RSA -alias <YOUR_ALIAS> -file
certreq.txt -keystore <PATH_TO_NEW_KEYSTORE_FILE> -storepass
<KEYSTORE_PASSWORD> -keypass <KEYSTORE_PASSWORD>
```

4. If you already have an SSL certificate, skip this step.

If you need an SSL certificate, use your CSR to obtain a new certificate from a third-party trusted SSL certificate provider. Save the certificate file on your server, outside the ScaleIO Gateway installation directory.

5. Import the Trusted Root, by typing this command. (If you did not define a keystore password, omit the password flags.)

```
keytool -import -alias root -keystore
<PATH_TO_NEW_KEYSTORE_FILE> -trustcacerts -file <LOCATION
OF_YOUR_root.cer_FILE> -storepass <KEYSTORE_PASSWORD> -
keypass <KEYSTORE_PASSWORD>
```

Note

The certificate must be in x.509 format.

If a message appears saying that the root is already in the system-wide store, import it anyway.

6. Import the intermediate certificates, by typing the command. (If you did not define a keystore password, omit the password flags.)

```
keytool -import -alias intermediateCA -keystore
<PATH_TO_NEW_KEYSTORE_FILE> -trustcacerts -file
<LOCATION_OF_YOUR_intermediate.cer_FILE> -storepass <keystore
password> -keypass <keystore password>
```

You must provide a unique alias name for every intermediate certificate that you upload with this step.

7. Install the SSL Certificate under the same alias that the CSR was created from (<YOUR_ALIAS> in previous steps), by typing the command (if you did not define a keystore password, omit the password flags):

```
keytool -import -alias <YOUR_ALIAS> -keystore
<PATH_TO_NEW_KEYSTORE_FILE> -trustcacerts -file
<LOCATION_OF_SSL_CERTIFICATE> -storepass <keystore password> -
keypass <keystore password>
```

8. Edit the following items in the file <ScaleIO_GATEWAY_INSTALLATION_DIRECTORY>\conf\catalina.properties:
 - a. keystore.file=<PATH_TO_NEW_KEYSTORE_FILE>
 - b. keystore.password=<PASSWORD_DEFINED_DURING_KEYSTORE_CREATION>

If you did not define a password, the default password is `changeit`.

9. Restart the scaleio-gateway service:
 - Windows
From the Windows **Services** management window, restart the EMC ScaleIO Gateway service.
 - Linux
Type the following command:

```
service scaleio-gateway restart
```

Replacement of the security certificate is complete.

Replacing the default self-signed security certificate with your own self-signed certificate

Procedure

1. Find the location of `keytool` on your server, and open it.

It is usually a part of the Java (JRE or JDK) installation on your server, in the `path` or `bin` directory. For example:

- a. C:\Program Files\Java\jdk1.7.0_25\bin\keytool.exe
- b. /usr/bin/keytool

2. Type the following command.

If you want to use the password generated by default (`changeit`), omit the `-keypass` and `-keystore` parameters from the command.

```
keytool -genkey -alias <YOUR_ALIAS> -keyalg RSA -validity 360
-keysize 2048 -storepass <KEYSTORE_PASSWORD> -keypass
<KEYSTORE_PASSWORD> -keystore <PATH_TO
_THE_CREATED_KEYSTORE_FILE>
```

Note

Store your keystore file in a directory outside the ScaleIO Gateway installation directory. This will prevent it from being overwritten when the ScaleIO Gateway is upgraded or reinstalled.

3. Edit the following items in the file `<ScaleIO_GATEWAY_INSTALLATION_DIRECTORY>\conf\catalina.properties`:

- a. `keystore.file=<PATH_TO_NEW_KEYSTORE_FILE>`
- b. `keystore.password=<PASSWORD_DEFINED_DURING_KEYSTORE_CREATION>`

If you did not define a password, the default password is `changeit`.

4. Restart the `scaleio-gateway` service:

- **Windows**
From the Windows **Services** management window, restart the EMC ScaleIO Gateway service.
- **Linux**
Type the following command:

```
service scaleio-gateway restart
```

Replacement of the security certificate is complete.

Upgrading the Gateway when a non-default certificate is used

If a non-default security certificate is used on the ScaleIO Gateway in the Windows and Linux environment (for example, if the certificate is signed by the user organization CA), 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
```

OpenStack interoperation with the ScaleIO Gateway

The OpenStack ScaleIO Cinder driver communicates with the ScaleIO Gateway through HTTPS (in other words, over SSL). By default, the driver ignores the gateway SSL certificate verification. However, the ScaleIO Cinder driver can be configured to verify the certificate.

Note

You can generate a self-signed certificate (.PEM file), using the keytool utility. For more information, see [“Generating a self-signed certificate using the keytool utility”](#).

To enable certificate verification, add the following parameters to the file `/etc/cinder/cinder_scaleio.config` on the Cinder node:

```
verify_server_certificate=true
server_certificate_path=<PATH_TO_PEM_FILE>
```

Generating a self-signed certificate using the keytool utility

This section describes how to generate self-signed certificates, using the keytool utility. The certificates can be used by the OpenStack ScaleIO driver, to communicate with the ScaleIO Gateway. For more information about configuring the driver, see [“OpenStack interoperation with the ScaleIO Gateway”](#).

To generate a self-signed certificate using the keytool utility, perform the following steps:

Procedure

1. Create a keystore file (.JKS) by typing the command:

```
keytool -genkeypair -keysize 1024 -alias herong_key -keypass
keypass -keystore herong.jks -storepass jkspass
```

2. Export the keystore file to a .PEM file, by typing the command:

```
keytool -exportcert -alias herong_key -keypass keypass -
keystore herong.jks -storepass jkspass -rfc -file
keytool_crt.pem
```

The certificate is stored in the file `<keytool_crt.pem>`. During configuration of the Cinder driver, the path to this .PEM file is required.

Advanced Installation Manager topics

This section describes advanced issues for Installation Manager installations, as well as how to use the IM REST API for installing.

Adding devices to SDS nodes on Windows servers

You can add devices to SDS nodes on Windows server by the following methods:

- Using a drive letter
- Using a file for the SDS with the `prepare_disk.exe` utility. This utility creates a mount point, formats it, and creates a file on it. To add the SDS device, use the full path of this file for the device path, as described in the following steps.

Using a physical disk or using a mount point on a raw partition are both not supported.

Note

As an alternative to manually running the `prepare_disk.exe` utility, you can use the CSV file during installation or extending of ScaleIO. To do so, in the **SDS Device List** column, add both drive letters and `PhysicalDriveX` (where *X* is the disk number from the Windows Disk Manager). You can also use this method with the CSV in situations where you need to add more devices than the amount of available drive letters for raw devices (d-z, excluding previously-allocated drives).

The `physicaldriveX` format is acceptable in the **SDS Device List** column of the CSV file. It is not case-sensitive.

Procedure

1. On the SDS, open a command line.
2. Verify that the file `prepare_disk.exe` exists on the SDS (folder location `C:\Program Files\EMC\scaleio\sds\bin`).
3. Run the command `prepare_disk.exe \\?\PhysicalDriveX`

where *X* is the disk number from the Windows Disk Manager.

Example:

```
C:\Program Files\EMC\scaleio> "C:\Program Files\EMC\scaleio
\sds\bin\prepare_disk.exe" \\?\PhysicalDrive8
```

Output similar to the following is displayed:

```
mosShmClient_Thrd:00481: Successfully connected to shared
memory server.
mosIO_close:00159: Closing FD(576)
```

4. Verify that the disk drive exists in the location `SYSTEMDRIVE\scaleio_devices`.

Example:

```
C:\scaleio_devices\PhysicalDrive8\PhysicalDrive8.bin
```

5. Add the new SDS device (one device per `scli --add_sds_device` command) using SCLI, or using the GUI:

Example:

```
scli --add_sds_device --sds_ip 10.76.61.6 --device_path C:\scaleio_devices\PhysicalDrive8\PhysicalDrive8.bin
```

```
C:\Program Files\EMC\scaleio\mdm\bin>scli --add_sds_device --sds_ip 10.76.61.6 --device_path C:\scaleio_devices\PhysicalDrive8\PhysicalDrive8.bin
```

Output similar to the following is displayed:

```
Successfully added device C:\scaleio_devices\PhysicalDrive8\PhysicalDrive8.bin to SDS. New device ID: 5b9f159300010002
```

Installing without validating Linux devices

During the Installation Manager validate stage, errors in the entering of Linux device names (such as `/dev/sdbb` instead of `/dev/sdb`) will prevent the continuation of the installation. This can be a problem if you want to add valid devices that do not meet the `/proc/partitions` standard (such as aliases and logical names).

To prevent the installation from getting stuck, you can select **Skip Linux devices validation** from the **Advanced settings** in the IM, as described in [step 4](#).

When this option has been selected, the validate stage will complete, but the affected node will show a completed (warning) status. You can click **Details** to see the error, and you can choose to proceed with the installation.

Using 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. The public key is placed on all servers that will allow access to the ScaleIO Gateway. The ScaleIO Gateway is the owner of the matching private key.

SioGWTool is used to upload the private key to the ScaleIO Gateway. SioGWTool is located in:

- **Linux:** `/opt/emc/scaleio/gateway/bin/SioGWTool.sh`
- **Windows:** `C:\Program Files\EMC\ScaleIO\Gateway\bin\SioGWTool.bat`

Note

In a new installation, if SSH key authentication will be used, there is no need to add node passwords to the CSV file used for system deployment.

Procedure

1. Create the private and public key pair:
 - a. Type the following command in the command line: `cd ~/.ssh/`
 - b. Type the following command: `ssh-keygen -t rsa`

This command generates a private and public key. When performing this command, you can add a passphrase, or generate the key without a passphrase. Two SSH keys are produced: `id_rsa` and `id_rsa.pub`.

2. On each ScaleIO server, type this command to store the public key in the `authorized_keys` file: `cat ~/.ssh/id_rsa.pub | cat >> ~/.ssh/authorized_keys`
3. Use SioGWTool in command line to upload the private key to the gateway, by typing the following command:

```
SioGWTool --update_ssh_key --ssh_key <PATH_TO_SSH_KEY_FILES>
[--key_passphrase <KEY_PASSPHRASE>]
```

The key passphrase is optional. Add this if you generated it when you created the key pair. During deployment, the gateway will use the private key for authentication.

Using the IM REST API

You can use the IM REST API to deploy a system. The IM is an orchestration engine that runs commands from different queues (nodes and MDM) at different "phases", and labels each one as a "process".

The IM REST API URLs begin with the `/im/` prefix.

This section includes the following topics:

General information

This section describes general information for the ScaleIO REST API, including REST API version in header fields, URIs, HTTP authentication, response fields, and ScaleIO Gateway configuration.

REST API version

The REST API version should be added to the `Accept` parameter in the header field. For example: `Accept: application/json;version=1.0`

If no `Accept` header field is present, it is assumed that the client accepts all media types and versions.

HTTP authentication

To perform authentication, invoke the HTTP GET request, with the URI: `/api/login`, with MDM `user\password` passes in HTTP Basic authentication. A token is returned.

For example:

```
https://10.76.60.190:443/api/login
```

RESPONSE

200 OK

HEADERS

Cache-Control:	no-cache	form	BODY
Cache-Control:	no-store		"YwRtaW46MTQwODY0ODcxODYwMTpmYzRmZDg1MTFmYmNiOGU5YzQ5ZTFmOGZjNTIxOTcxYw"
Content-Encoding:	gzip		
Content-Type:	application/json; charset=UTF-8		
Date:	2014 Aug 21 14:18:38		download
Expires:	1970 Jan 1 02:00:00 -44 years		
Pragma:	no-cache		
Server:	Apache-Coyote/1.1		
Transfer-Encoding:	chunked		
Vary:	Accept-Encoding		

This token is valid for 8 hours from the time it was created, unless there has been no activity for 10 minutes, or if the client has sent a logout request.

HTTP token invalidation (logout)

The invalidation is done by invoking the HTTP GET request, with the URI: `/api/logout`

The token mentioned above is the password in the HTTP Basic authentication (the user is ignored - it can be empty).

For every REST API request that does not require the `gatewayAdminPassword`, the authentication is done by passing the token mentioned above as the password in the HTTP Basic authentication (the user is ignored - it can be empty).

Requests that require the `gatewayAdminPassword` work similarly, except that instead of `/api/login`, invoke an HTTP GET request, `/api/gatewayLogin` with user: admin password: `<gatewayAdminPassword>` in HTTP Basic authentication. A token is returned. Instead of invoking `/api/logout`, invoke `/api/gatewayLogout` with the token received when you logged in.

Note

Requests that require `gatewayAdminPassword` are:
GET:

```
/api/Configuration
/api/gatewayLogout
/api/getHostCertificate/{Mdm | Lia}
```

POST:

```
/api/updateConfiguration
/api/instances/System/action/createMdmCluster
/api/trustHostCertificate/{Mdm | Lia}
/api/gatewaySetSecureCommunication
```

Response fields

The order of the fields in the responses may change. More fields may be added in the future.

URIs

- POST (create) / GET all objects for a given type:
`/api/types/{type}/instances`
- GET by id:
`/api/instances/{type::id}`

- **POST a special action on an object:**
`/api/instances/{type::id}/action/{actionName}`
- **POST a special action on a given type:**
`/api/types/{type}/instances/action/{actionName}`
- **Get current API version:**
`/api/version`
- **Every row in the Object's Parent table appears as a link in the response of get object:**
`/api/instances/{type::id}`
- **Every row in the Object's Relationships table appears as a link in the response of get object:**
`/api/instances/{type::id}/relationships/{Relationship name}`
- **GET all instances**
`/api/instances/`

Table 17 Response

Property	Type	Note
sessionTag	Long	
isDirty	Boolean	The version on some objects changed while the MDM created the response
System	Syst	
lastSystemVersion	Long	
protectionDomainList	List of ProtectionDomain objects	
lastProtectionDomainVersion	Long	
sdsList	List of SDS objects	
lastSdsVersion	Long	
storagePoolList	List of StoragePool objects	
lastStoragePoolVersion	Long	
deviceList	List of Device objects	
lastDeviceVersion	Long	
volumeList	List of Volume objects	
lastVolumeVersion	Long	
vTreeList	List of VTree objects	
lastVTreeVersion	Long	
sdcList	List of SDC objects	

Table 17 Response (continued)

Property	Type	Note
lastSdcVersion	Long	
faultSetList	List of FaultSet objects	
lastFaultSetVersion	Long	
rfcacheDeviceList	List of RfcacheDevice objects	
rfcacheDeviceVersion	Long	

- **Change configuration of ScaleIO Gateway (POST)**

/api/updateConfiguration/

Request:

Parameters (AND\OR):

- `mdmAddresses`—List of MDM IP addresses that the ScaleIO Gateway will use to connect to the MDM (for performance improvement, place the master MDM IP addresses before the slave ones)
- `mdmPort`—MDM port
- `gatewayAdminPassword`—Password for installation manager, get/update ScaleIO gateway configuration, and add master MDM
- `systemId`—ID of the system configured by `mdmAddresses`. If `systemId` is left empty, it will be populated on the first invocation of `/api/login`. At every `/api/login`, the system ID value is compared to the ID of the ScaleIO cluster (MDM). If the IDs are different (`mdmAddresses` points to a cluster that does not match the `systemId`), an error message is displayed, and all login tokens (`/api/login`, `/api/gatewayLogin`) are invalidated. None of the gateway clients will be able to send requests to the MDM until one of the following happens:
 - `systemId` is configured to the correct value
 - `mdmAddresses` is configured to the correct IP addresses
 - `systemId` is set to an empty value. In this case, on the first invocation of `/api/login`, the value for `systemId` will be populated according to the ID of the matching `mdmAddresses`.

If the value for `systemId` is different from the configured `systemId` (and the configured `systemId` is not empty), all login tokens (`/api/login`, `/api/gatewayLogin`) are invalidated.
- `snmpSamplingFrequency`—MDM sampling frequency in seconds. If sampling frequency is set to 0, the SNMP trap sender will be disabled, and no sampling to the MDM will be performed.
- `snmpResendFrequency`—Resend frequency of SNMP traps, in minutes. If resend frequency is set to 0, all traps will be sent in every sampling.
- `snmpTrapsReceiverIp`—SNMP trap receivers' IP addresses (supports up to two comma-separated or semi-colon-separated IP addresses)
- `snmpPort`—The port number used for SNMP traps

For example:

```
{
  "mdmAddresses":["10.76.60.150", "10.76.60.11"],
  "mdmPort":"6611",
  "gatewayAdminPassword":"Password1",
  "systemId":"7f5d8fc72a3d7f3d" ,
  "snmpSamplingFrequency":"30",
  "snmpResendFrequency":"0",
  "snmpTrapsReceiverIps":["10.76.60.190","10.76.60.191"],
  "snmpPort":"162"
}
```

- **Get configuration of ScaleIO Gateway (GET)**

/api/Configuration/

Response:

- **mdmAddresses**—List of MDM master IP addresses that the ScaleIO Gateway will use to connect to the MDM
- **mdmPort**—MDM port
- **mdmUsername**—MDM user name
- **systemId**—ID of the system, configured by **mdmAddresses** (see detailed explanation above, for "Change configuration of ScaleIO Gateway (POST)".)
- **snmpSamplingFrequency**—MDM sampling frequency in seconds
- **snmpResendFrequency**—Resend frequency of SNMP traps, in minutes
- **snmpTrapsReceiverIp**—SNMP trap receivers' IP addresses
- **snmpPort**—The port number used for SNMP traps
- **remoteSyslog**—The remote syslog servers configuration in the MDM
- **featuresEnableSnmp**—Indicates whether SNMP is enabled or not
- **cipherSuites**—A list of names of cipher suites (as recognized by Java) that the SSL will use instead of the default list.
- **featuresEnableIM**—Indicates whether the Installation Manager is enabled or not
- **allowNonSecureCommunication**

For example:

```
{
  "snmpResendFrequency": "0",
  "snmpSamplingFrequency": "30",
  "snmpPort": "162",
  "mdmPort": "6611",
  "remoteSyslog": [
    {
      "hostName": "10.76.60.100",
      "port": 1468,
      "facility": 16
    },
    {
      "hostName": "10.76.60.101",
      "port": 1468,
      "facility": 16
    }
  ],
  "mdmAddresses": [
```

```

"10.76.60.150",
"10.76.60.135"
],
"systemId":"7f5d8fc72a3d7f3d",
"snmpTrapsReceiverIps": [
    "10.76.60.192"
]
}

```

- **Query selected statistics (POST)**

/api/instances/querySelectedStatistics

Request:

- **Required parameters:**

`selectedStatisticsList`—list of objects containing:

- `type`—object type (System, ProtectionDomain, Sds, StoragePool, Device, Volume, VTree, Sdc, FaultSet, RfcacheDevice)
- `properties`—list of properties to fetch

- The following parameters are not relevant to the System type (can be omitted):

`ids`—list of objects ids

or

`allIds`—with empty value

Response:

The response contains a list of all requested types, the ids, and properties, according to the request parameters.

The System type does not contain an id.

Types for which the requested ids do not exist in the cluster will appear with an empty {} - see `FaultSet` in the example below.

Any id (from the `ids` parameters in the request) that does not exist in the cluster will not appear in the response. The following will appear in the `scaleio.log`:

```
Got no statistics for <type>::<id>.
```

In the following example, `Sds::022beb23000000004` does not exist in the cluster, so it does not appear in the response. Instead, this will appear in the `scaleio.log`:

```
Got no statistics for Sds::id: 022beb23000000004
```

Example:

POST `https://localhost:8443/api/instances/querySelectedStatistics`

Body:

```

{"selectedStatisticsList": [
  {"type": "ProtectionDomain", "ids": ["cc480c9b00000000"],
  "properties": ["capacityInUseInKb">,
    {"type": "Volume", "ids":
    ["022beb25000000006", "022beb23000000004"], "properties":
    ["numOfMappedSdcs", "userDataWriteBwc">,
    {"type": "Sds", "ids": ["c919d820000000001", "022beb23000000004"],

```

```

"properties":["capacityInUseInKb">,
  {"type":"System", "allIds":"", "properties":
  ["rmcacheSizeInKb">,
    {"type":"FaultSet", "ids":
    ["c919d82000000001","022beb2300000004"], "properties":
    ["numOfSds">,
      {"type":"StoragePool", "allIds":"", "properties":
      ["unreachableUnusedCapacityInKb", "numOfThinBaseVolumes">
        >

```

The response:

```

{
  "FaultSet":{
  },
  "Volume":{
    "022beb2300000004":{
      "userDataWriteBwc":{
        "numOccured":0,
        "totalWeightInKb":0,
        "numSeconds":1
      },
      "numOfMappedSdcs":0
    },
    "022beb2500000006":{
      "userDataWriteBwc":{
        "numOccured":0,
        "totalWeightInKb":0,
        "numSeconds":1
      },
      "numOfMappedSdcs":0
    }
  },
  "StoragePool":{
    "7db925c200000003":{
      "numOfThinBaseVolumes":0,
      "unreachableUnusedCapacityInKb":0
    },
    "7db925bf00000000":{
      "numOfThinBaseVolumes":2,
      "unreachableUnusedCapacityInKb":0
    },
    "7db925c000000001":{
      "numOfThinBaseVolumes":0,
      "unreachableUnusedCapacityInKb":0
    },
    "7db925c100000002":{
      "numOfThinBaseVolumes":0,
      "unreachableUnusedCapacityInKb":0
    }
  },
  "ProtectionDomain":{
    "cc480c9b00000000":{
      "capacityInUseInKb":50331648
    }
  },
  "System":{
    "rmcacheSizeInKb":524288
  },
  "Sds":{
    "c919d82000000001":{
      "capacityInUseInKb":12509184
    }
  }
}

```

- Get MDM's or LIA's certificate (GET)

```
/api/getHostCertificate/{Mdm|Lia}?host={host ip}
```

The request can be sent for either MDM or LIA.

Example:

```
/api/getHostCertificate/Mdm?host=10.76.60.10
```

Response:

The host certificate in PEM encoding.

Note

The whole certificate should be saved into a `cer` file for the `trustHostCertificate` request.

For example:

```
-----BEGIN CERTIFICATE-----
MIIDcTCCAlmgAwIBAAIBATANBgkqhkiG9w0BAQUFADB8MQwwCgYDVQQqEwNNRE0xFzA
VBgNVBAMTDmNlbnRvcy02LTQtYWRpMRIwEAYDVQQHEw1Ib3BraW50b24xZjAUBgNVBA
gTDU1hc3NhY2h1c2V0dHMxCzAJBgNVBAYTA1VTMQwwCgYDVQQKEwNFTUMxDDAKBgNVB
AsTA0FTRDAeFw0xNTEyMTkwNzI4MTVaFw0yNTEyMTcwODI4MTVaMHwxDDAKBgNVBCoT
A01ETTEXMBUGA1UEAxMOY2VudG9zLTltYtNC1hZGkxEjAQBgNVBAcTCUhhvGtpbnRvbGJ
WMBQGA1UECBMNTWFzc2FjaHVzZXR0czELMAkGA1UEBhMCVVMxDDAKBgNVBAoTA0VNQz
EMMAoGA1UECXMDOVNEMIIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAA4SMyb
aAEZjfBX9wLglr3wxYHOvID5Pe1Z26Pv8oIR/
MTOVa1Bw4A9px1MHHSIfkAfgRlLC24uebZXhbb0snBq+OL+SJPwEfbOVbif/
saXL8RJFwm/VNg8KHUwjuq/sJkKDjx9uSf0U+/9FzwwKVuM87xDj/
rVvJgBYh6pH34q/XD5l8am/iEQr/EnGZmIsa+VkcL0IeYKbkA3ZINfI4YsJSJ
+qeu5e/
KMsNlHEvmhk1DdJbLayn9QkiS5Q9e8A40jjkb2e1Q71awo0lb6+8XXWWkpBhxAnRa9P
8Pb1BfcNyUfXtrKuy+fRjw4Gp
+rw2MdoIDuMb0l+1sQaRvVPTYxwIDAQABMA0GCSqGSIB3DQEBBQUAA4IBAQCule/
jBz63ZFDs+pFvZ3XI/VMdn9NArM
+8Cjn7Luar8oEVAYI6jqYYcZCk2jQyfuI1HP2jXqPTJo8CNyNT5S6rZ5ryHOiRjn/
K2pVC6kT497LY5lc3LjhXUdLjpWnW2jsGfM93cCkkrxu8wmkh9oo8WizOiRAyK mz02u
TEuEok7GJBS/
DR6csnLo2YLUV6ZqeBN9jdzZbIY7SoFWya1K4xZmqhkAtnj1ynP3uoxTkd
+wfdRmYeDv8l5eciLj2BXNuV8zXYWSCyABZC//jvajNtSEXgUura3uh0YBIfbO/
AZ980zUMwJBMBR06yw4tHnHRRYgfI3tnZOD4byaJODHuq
-----END CERTIFICATE-----
```

- **Trust an MDM's or LIA's certificate (POST)**

```
/api/trustHostCertificate/{Mdm|Lia}
```

Example:

```
/api/trustHostCertificate/Mdm
```

Request:

Content-type: multipart/form-data

The body should contain a part named "file" and the file containing the certificate to be trusted.

- **Set the gateway to work with secured communication with the MDM (POST):**

```
/api/gatewaySetSecureCommunication
```

The gateway will not be able to connect to the MDM using non-secured communication.

- Working with the Installation Manager (IM) REST API

IM is an orchestration engine that runs commands from different queues (nodes and MDM) at different "phases," and labels them as a "process" (such as an upgrade).

The IM REST API begins with the `/im/` prefix.

Login to the IM REST API must be done via POST to the `j_spring_security_check` url, followed by username, password, and the login submission. For example:

```
POST https://localhost/j_spring_security_check
Content: "j_username=admin&j_password=Password1&submit=Login"
```

An operation such as the above will provide a `JSESSIONID` cookie in the response that should be used for identification of the session in all the future requests.

For example:

```
POST "Content-Type: application/json" "Cookie:
JSESSIONID=969F624A761937AE80E6CC9E91756B10" https://
localhost/im/types/Command/instances/actions/retry
```

- Most IM operations require a topology / configuration object as a parameter (payload) for the REST (HTTP) request in JSON format.

Such an object has the following format:

Note

Each field's functionality is marked with the `/* */` notation—it should not be used when generating such instance with JSON format {

`/* target IO address to send SNMP traps while configuring Gateway to become SNMP sampler */`

```
"snmpIp": null,
```

`/* list of IP addresses from which LIA will only accept connections—if null, accept any connection */`

```
"safeIPsForLia": null,
```

`/* system installation ID to be used or configured*/`

```
"installationId": null,
```

`/* all MDM manager IP addresses list*/`

```
"mdmIPs": ["10.76.60.48",
"10.76.60.41"],
```

/* password for MDM and LIA */

```
"mdmPassword": "Password1",
"liaPassword": "Password1",
```

/* deprecated—not in use */

```
"licenseKey": null,
```

/* deprecated—not in use */

```
"licenseType": "unlimitedCapacity",
```

/* boolean, indicating whether cluster should be set with high performance profile */

```
"isClusterOptimized": null,
```

/* deprecated—not in use */

```
"callHomeConfiguration": null,
```

/* string describing parameters for remote sys log configuration */

```
"remoteSyslogConfiguration": null,
```

/* string describing cluster (master) version */

```
"systemVersionName": "",
```

/* boolean indicating whether system is running */

```
"upgradeRunning": false,
```

/* boolean indicating whether current MDM is a cluster (if running) */

```
"clustered": false,
"masterMdm": {
  "node": {
    "ostype": "linux",
    "nodeName": null,
    "nodeIPs": ["10.76.60.41"],
```

/* domain only relevant (if at all) for Windows */

```
"domain": null,
"userName": "root",
"password": "Password1",
```

```
        "liaPassword": null
    },
```

/* internal data fields—do not use */

```
    "nodeInfo": null,
```

/* in case of upgrade—if rollback should occur—to which version */

```
    "rollbackVersion": null,
    "mdmIPs": ["10.76.60.41"],
    "name": "M2",
```

/* unique MDM ID */

```
    "id": null,
```

/* internal data fields—do not use */

```
    "ipForActor": null,
```

/* IP addresses to be used and configured for management */

```
    "managementIPs": ["10.76.60.41"]
},
```

/* similar to master MDM */

```
    "slaveMdmSet": [{
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.48"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
            "liaPassword": null
        },
        "nodeInfo": null,
        "rollbackVersion": null,
        "mdmIPs": ["10.76.60.48"],
        "name": "M1",
        "id": null,
        "ipForActor": null,
        "managementIPs": ["10.76.60.48"]
    }],
```

/* similar to master MDM */

```
    "tbSet": [{
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.146"],
            "domain": null,
```

```

        "userName": "root",
        "password": "Password1",
        "liaPassword": null
    },
    "nodeInfo": null,
    "rollbackVersion": null,
    "mdmIPs": ["10.76.60.146"],
    "name": "M3",
    "id": null,
    "tbIPs": ["10.76.60.146"]
}],

```

/* similar to master MDM */

```

"standbyMdmSet": [],

```

/* similar to master MDM */

```

"standbyTbSet": [],
"sdsList": [{
    "node": {
        "ostype": "linux",
        "nodeName": null,
        "nodeIPs": ["10.76.60.146"],
        "domain": null,
        "userName": "root",
        "password": "Password1",
        "liaPassword": null
    },
    "nodeInfo": null,
    "rollbackVersion": null,
    "sdsName": "SDS_10.76.60.146",

```

/* name (required) and id (optional, if it already exists) for Protection Domain */

```

"protectionDomain": "domain1",
"protectionDomainId": null,

```

/* fault set id (if defined) */

```

"faultSet": null,

```

/* various SDS IP addresses and their usage */

```

"allIPs": ["10.76.60.146"],
"sdsOnlyIPs": null,
"sdcOnlyIPs": null,
"devices": [],

```

/* Read Flash Cache definition if defined, on which pools and which devices */

```

"rfCached": false,
"rfCachedPools": [],
"rfCachedDevices": [],

```

`/* boolean, indicating whether SDS is set to high performance profile */`

```

    "optimized": false,
    "port": 7072,
    "id": "0"
  },
  {
    "node": {
      "ostype": "linux",
      "nodeName": null,
      "nodeIPs": ["10.76.60.48"],
      "domain": null,
      "userName": "root",
      "password": "Password1",
      "liaPassword": null
    },
    "nodeInfo": null,
    "rollbackVersion": null,
    "sdsName": "SDS_10.76.60.48",
    "protectionDomain": "domain1",
    "protectionDomainId": null,
    "faultSet": null,
    "allIPs": ["10.76.60.48"],
    "sdsOnlyIPs": null,
    "sdcOnlyIPs": null,
    "devices": [],
    "rfCached": false,
    "rfCachedPools": [],
    "rfCachedDevices": [],
    "optimized": false,
    "port": 7072,
    "id": "0"
  },
  {
    "node": {
      "ostype": "linux",
      "nodeName": null,
      "nodeIPs": ["10.76.60.41"],
      "domain": null,
      "userName": "root",
      "password": "Password1",
      "liaPassword": null
    },
    "nodeInfo": null,
    "rollbackVersion": null,
    "sdsName": "SDS_10.76.60.41",
    "protectionDomain": "domain1",
    "protectionDomainId": null,
    "faultSet": null,
    "allIPs": ["10.76.60.41"],
    "sdsOnlyIPs": null,
    "sdcOnlyIPs": null,
    "devices": [],
    "rfCached": false,
    "rfCachedPools": [],
    "rfCachedDevices": [],
    "optimized": false,
    "port": 7072,
    "id": "0"
  }
],
"sdcList": [{
  "node": {
    "ostype": "linux",
    "nodeName": null,
    "nodeIPs": ["10.76.60.41"],
    "domain": null,
    "userName": "root",
    "password": "Password1",

```

```
        "liaPassword": null
    },
    "nodeInfo": null,
    "rollbackVersion": null,
```

/* splitter (RecoverPoint) IP address to replicate the IO */

```
"splitterRpaIp": null,
```

/* Boolean indicating whether SDC should be set to high performance profile */

```
        "optimized": false
    },
    {
        "node": {
            "ostype": "linux",
            "nodeName": null,
            "nodeIPs": ["10.76.60.48"],
            "domain": null,
            "userName": "root",
            "password": "Password1",
            "liaPassword": null
        },
        "nodeInfo": null,
        "rollbackVersion": null,
        "splitterRpaIp": null,
        "optimized": false
    }
}],
```

/* security setting—whether to allow process to connect to non-secured LIA / MDM, and whether ScaleIO components should be configured to be non-secured (secured by default) */

```
"securityConfiguration": {
    "allowNonSecureCommunicationWithMdm": false,
    "allowNonSecureCommunicationWithLia": false,
    "disableNonMgmtComponentsAuth": false
},
}
```

- The Node object is used as an attribute of each ScaleIO component (for example, on which node the MDM is installed), and describes how the node will be accessed:
 - via SSH (Linux node with root credentials)
 - via WMI (Windows node with admin domain and credentials)
 - via LIA (unknown node with null domain, null credentials and LIA password)

For example:

```
//linux node accessed via SSH
"node": {
    "ostype": "unknown",
    "nodeName": null,
    "nodeIPs": ["10.76.60.48"],
    "domain": null,
    "userName": null,
    "password": null,
    "liaPassword": null
```

```

    },
    //unknown node (linux or windows) - accessed via LIA
    "node": {
        "ostype": "linux",
        "nodeName": null,
        "nodeIPs": ["10.76.60.41"],
        "domain": null,
        "userName": "root",
        "password": "Password1",
        "liaPassword": null
    },

```

- **LIA usage**—Any post installation operations **MUST** be done via LIA. Such operations (for example, upgrades) may fail unless all nodes are targeted to be performed via LIA.

Linux node accessed via SSH

```

"node": {
  "ostype": "unknown",
  "nodeName": null,
  "nodeIPs": ["10.76.60.48"],
  "domain": null,
  "userName": null,
  "password": null,
  "liaPassword": null
},

```

Unknown node (Linux or Windows) - accessed via LIA

```

"node": {
  "ostype": "linux",
  "nodeName": null,
  "nodeIPs": ["10.76.60.41"],
  "domain": null,
  "userName": "root",
  "password": "Password1",
  "liaPassword": null
},

```

How to use IM REST API URIs

JSESSION ID must be used to in order to run IM REST API URIs. URIs run without JSESSION ID will fail.

The following steps can be used to obtain a JSESSION ID.

Procedure

1. Run the following URI:

```

https://localhost/j_spring_security_check
Content: "j_username=admin&j_password=Password1&submit=Login"
POST

```

where the user name in the above example is "admin" and the password is "Password1".

2. Find the JSESSION ID in the URI output. For example:

```
JSESSIONID=969F624A761937AE80E6CC9E91756B10
```

- Set the following headers, using the JSESSION ID obtained in the previous steps. For example:

```
"Content-Type": "application/json"
"Cookie": "JSESSIONID=969F624A761937AE80E6CC9E91756B10"
```

Results

The JSESSION ID will be added to the URIs. You may now run any IM REST API URI.
Figure 18 Example of IM REST API URI when JSESSION ID is configured as a header

The screenshot displays a REST client interface with a dark header bar containing the logo and tabs for 'REQUESTS' and 'SCENARIOS'. Below the header, there is a search bar with the placeholder text 'type a name'. The main area is divided into two sections: 'REQUEST' and 'RESPONSE'.

REQUEST Section:

- Method:** HTTPS (indicated by a lock icon and a dropdown menu).
- URL:** `10.76.61.8/im/types/ProcessPhase/actions/moveToldlePhase`
- HEADERS:**
 - ☒ Content-Type: `application/json`
 - ☒ Cookie: `JSESSIONID=2340E4`
- BODY:** `1 {}`

RESPONSE Section:

- Status:** 200 OK (displayed in a green bar).
- HEADERS:**
 - Cache-Control: `no-cache`
 - Cache-Control: `no-store`
 - Content-Encoding: `gzip`
 - Date: `2016 Sep 11 15:40:30 -2h 8m`
- BODY:** (The response body is obscured by a large 'NO CO' watermark).

Deploying a system

When using the IM REST API to deploy a system, use JSON objects instead of the CSV file, and the appropriate steps of IM REST API commands to perform these activities:

Procedure

1. Upload the software packages
2. Install the software packages
3. Configure the system
4. Monitor progress

Any post installation operations **MUST** be performed via LIA. Operations such as upgrades may fail unless all nodes have been targeted to be accessed via LIA.

To deploy the system, use the following work flow. A detailed example of a JSON configuration is shown in the *REST API Reference Guide*.

5. Upload packages to the Installation Manager:

```
POST /im/types/InstallationPackages/instances/actions/
uploadPackages
```

6. Install the packages (query phase starts automatically), using the JSON configuration:

```
/im/types/Configuration/actions/install
```

7. Check progress until the installation phase is finished:

```
/im/types/Command/instances
```

8. When the installation phase is completed, move to the next phase:

```
/im/types/ProcessPhase/actions/moveToNextPhase
```

Advanced vSphere plug-in topics

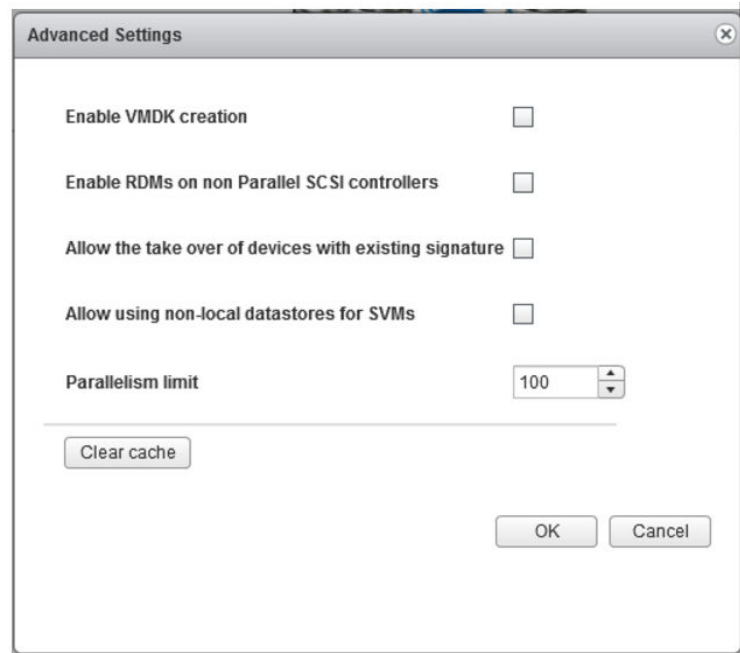
This section describes advanced issues for plug-in installation.

Advanced settings options

This topic describes the advanced settings options that you can use to configure the VMware plug-in deployment settings.

To access these settings, click from the **EMC ScaleIO** screen, click **Advanced settings**.

The **Advanced Settings** dialog appears:



Enable the features you want before starting the ScaleIO deployment, or, if you are in the midst of a deployment, you can click **Cancel**, then make the changes. When you restart the deployment, you can pick up from where you left off.

- **Enable VMDK creation**
Enables the addition of devices as VMDK, as opposed to RDM.
- **Enable RDMS on non Parallel SCSI controllers**
Enables non-SCSI controEnable RDMS on non Parallel SCSI controllerssler devices to be added as RDM. Note: Do not enable this option if the device does not support SCSI Inquiry Vital Data Product (VPD) page code 0x83.
- **Allow the take over of devices with existing signature**
Enables the reuse of devices that were part of a previous ScaleIO system, and were not removed from that system properly. This will revert to not selected after use, and must be reactivated.
- **Allow using non-local datastores for ScaleIO gateway**
Enable deployment of the Gateway SVM on non-local datastores (such as remote storage), in case of insufficient space on the local DS, or if no DS can be found.
- **Parallelism limit**
Enables the increase of the parallelism limit (default: 100), thus speeding up the deployment, which can be useful in deployment of a very large ScaleIO system (several hundred nodes). This is dependent on the processing power of the vCenter.

Registering the ScaleIO plug-in manually

This topic describes an advanced way to use PowerCLI to register the ScaleIO plug-in on a vCenter.

The plug-in is provided as a ZIP file that can be downloaded by the vSphere web client servers in your environment. The ZIP file can be downloaded directly from the EMC Online Support site, or, if the web servers may not have internet access, from a file server.

Before you begin, ensure that there is communication between the vSphere web client server and the web server storing the plug-in.

Procedure

1. You can upload the ZIP file to an HTTP or an HTTPS server.

If you are uploading the ZIP file to an HTTP server, perform the following:

- a. On the computer where the vSphere Web client is installed, locate the `webclient.properties` file.

vCenter	Operating system	Path to file
5.x	Windows 2003	%ALLUSERSPROFILE%\Application Data\VMware\vSphere Web Client
	Windows 2008/2012	%ALLUSERSPROFILE%\VMware\vSphere Web Client
	Linux	/var/lib/vmware/vsphere-client
6.x	Windows 2008/2012	C:\ProgramData\VMware\vCenterServer\cfg\vsphere-client
	Linux	/etc/vmware/vsphere-client/

- b. Add the following line to the file:

```
allowHttp=true
```

- c. Restart the VMware vSphere Web Client service.

2. Using PowerCLI for VMware, and set to **Run as administrator**, run `Set-ExecutionPolicy AllSigned`.
3. Close PowerCLI, and reopen it, set as **Run as administrator**.
4. Extract `EMC-ScaleIO-vSphere-plugin-installer-2.5-<build>.X.zip`
5. Upload `EMC-ScaleIO-vSphere-web-plugin-2.5-<build>.X.zip` to the HTTP/HTTPS server.
6. From PowerCLI, run the following script: `ScaleIOPluginSetup-2.5-<build>0.X.ps1`
 - a. Enter the vCenter name or IP address, user name, and password.
 - b. For **Select Mode**, choose option 1, `Register SIO plugin`.
 - c. For **Select Registration Mode**, choose `Advanced`.
 - d. Enter the full URL of the plug-in ZIP file.
 If the ZIP file is located on a ScaleIO Gateway, enter the following URL:
`https://<SCALEIO_GATEWAY_IP_ADDRESS>`.
 If you are manually placing the zip file on the server, place it in the `/opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository` folder.
 - e. If necessary, accept the thumbprint.
7. Log out, and log back in to the vSphere web client.

Example for HTTP server:

```
.\registerScaleIOPlugin.ps1 -vcenter 10.103.109.16
    -userUrl "http://10.76.60.14/sample/ScaleIO-vSphere-web-
plugin-1.30.0.160.zip"
```

Example for HTTPS server:

```
.\registerScaleIOPlugin.ps1 -vcenter 10.103.109.16
```

```
-userUrl
"https://10.76.61.139/sample/ScaleIO-vSphere-web-
plugin-1.30.0.160.zip"
-thumbprint CA:66:49:D0:CE:D9:8C:A0:D0:93:E3:83:DE:59:25:5F:
79:E1:53:B6
-adminEmail test.email@emc.com
```

The script registers the plug-in and the following message appears:

```
Registering ScaleIO extension...
```

```
The extension was registered successfully
```

Troubleshooting plug-in registration issues

You can use the following logs to assist in troubleshooting problems that may occur during registration of the VMware plug-in. To find relevant log entries, search for `scaleio` in the log file.

The vSphere web client (Virgo) logs are located in the following directories:

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

Other relevant logs:

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

Deploy ScaleIO in a 2-layer environment

Deploy the ScaleIO MDM and SDS components on physical Linux servers, and the SDCs on ESXi servers.

You can deploy ScaleIO in a 2-layer environment using one of the following methods:

- [Deploying ScaleIO in a 2-layer environment using the Installation Manager and plug-in](#) on page 217
- [Manually deploying ScaleIO in a 2-layer environment](#) on page 218

Deploying ScaleIO in a 2-layer environment using the Installation Manager and plug-in

Deploy the ScaleIO MDM and SDS components on physical Linux servers, and the SDCs on ESXi servers.

The steps in this task are described fully in the *ScaleIO Deployment Guide*.

Procedure

1. Deploy the physical Linux environment:
 - a. Install the ScaleIO Gateway, as described in [Preparing the Installation Manager and the Gateway](#) on page 68.
 - b. Deploy ScaleIO, as described in [Installing with the full Installation Manager](#) on page 72.

Note

When preparing the CSV, you must enter No in the Is SDC column for all nodes.

 - c. Install the ScaleIO GUI, as described in [Install the ScaleIO GUI](#) on page 102.
2. Register the Linux system:
 - a. Register the ScaleIO plug-in, as described in [Registering the ScaleIO plug-in](#) on page 109.
Do not upload the OVA template.
 - b. Log in to the vSphere client.
 - c. From the ScaleIO **Basic tasks** list, click **Register ScaleIO system**.
 - d. Follow the steps to register the Linux system.
3. Install the SDC on ESX, as described in [Preparing the ESXi hosts](#) on page 113.
4. Deploy ScaleIO:
 - a. Start the deployment, as described in [Deploying ScaleIO with RDM/VMDK device management](#) on page 124.
 - b. In the **Select Installation** screen, select **Add servers to a registered ScaleIO system**.
 - c. In the **Add ESX hosts to cluster** screen, select the relevant datacenter.
 - d. Click **Next** through the following screens until you reach the **Add SDCs** screen.

- e. In the **Add SDCs** screen:
 - a. Type the ESX root password.
 - b. Select **Disable** for the SCSI comparison.
- f. Click **Next** through the following screens until you reach the **Configure Upgrade Components** screen.
- g. Clear the **Install ScaleIO Gateway** option.
- h. In the **Configure Upgrade Components screen**, select the management and data networks.
- i. Review the deployment summary and click **Finish**.
- j. Enter the vCenter password.

The deployment process (adding the SDC) begins.

Manually deploying ScaleIO in a 2-layer environment

Manually deploy the ScaleIO MDM and SDS components on physical Linux servers, and the SDCs on ESXi servers.

The steps in this task are described fully in the *ScaleIO Deployment Guide*.

Procedure

1. Manually deploy the physical Linux environment:
 - a. Install the MDM, as described in [MDM cluster setup](#) on page 160.
You can set up the MDM cluster in 3-node, 5-node, or single-node mode.
 - b. Install and enable RfCache, if using, as described in [Installing RfCache - Linux](#) on page 172.
 - c. Install the ScaleIO GUI, as described in [Install the ScaleIO GUI](#) on page 102.
 - d. Install the Gateway, as described in [Installing the Gateway - Linux](#) on page 173.
 - e. Install the LIA, as described in [Installing the LIA - Linux](#) on page 174.
2. Register the Linux system:
 - a. Register the ScaleIO plug-in, as described in [Registering the ScaleIO plug-in](#) on page 109.
Do not upload the OVA template.
 - b. Log in to the vSphere client.
 - c. From the ScaleIO **Basic tasks** list, select **Register ScaleIO system**.
 - d. Follow the steps to register the Linux system.
3. Install the SDC on the ESX, as described in [Installing the SDC directly on an ESX host](#) on page 182.

You will need to reboot the ESX host during the installation.

Prepare the servers for deployment

Before deploying ScaleIO on ESXi servers, you must prepare the servers.

Before you begin

Ensure that you have:

- IP address of VMware vCenter server
- User name and password for accessing the vCenter with the vSphere Web Client

Before deploying ScaleIO systems, you must configure VMware security features. This procedure describes how to do so with the vSphere web client.

Note

Ensure that the remote connection is established within 1-2 minutes, otherwise, the lockdown mode is enabled on the ESXi servers, preventing you from subsequent remote connectivity. When in lockdown mode, you can connect to the ESXi servers locally.

Procedure

1. Log in to the vSphere web client.

The **vSphere Web Client** screen is displayed with **Navigator** and **Home** panes; the **Navigator** pane with network entities as nodes and the **Home** pane with network entities as icons.

2. In the **Navigator** pane, select **Home > Hosts and Clusters**.

Preconfigured hosts, clusters, and datacenters are displayed as navigation tree nodes, in the **Navigator** pane.

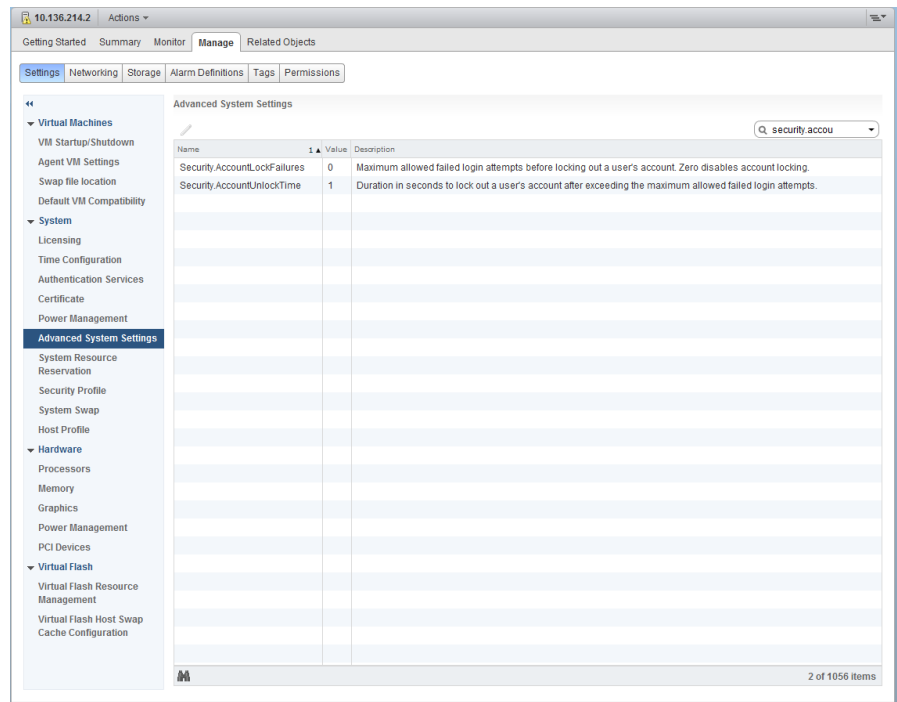
3. Configure the security-related parameters:

- a. Select an ESXi server, under a host/cluster/datacenter node.

The ESXi servers are built under a host, cluster, or datacenter node, in the **Navigator** pane.

- b. Under **Manage** menu, select **Settings > System > Advanced System Settings**.

A list of system-defined parameters, with predefined parameter values, is displayed in the work area.



c. Using the **Search** (🔍) field, filter the following security parameters:

- Set `Security.AccountLockFailures = 0`
The parameter value defines the maximum number of failed login attempts allowed, before locking out the user's account. The zero (0) value disables account locking.

Note

To disable the locking out of user accounts, the parameter value must be set as zero (0).

- Set `Security.AccountUnlockTime = 1 (in seconds)`
The parameter value defines the duration in seconds to lock out a user's account, after exceeding the maximum number of allowed failed login attempts.

Note

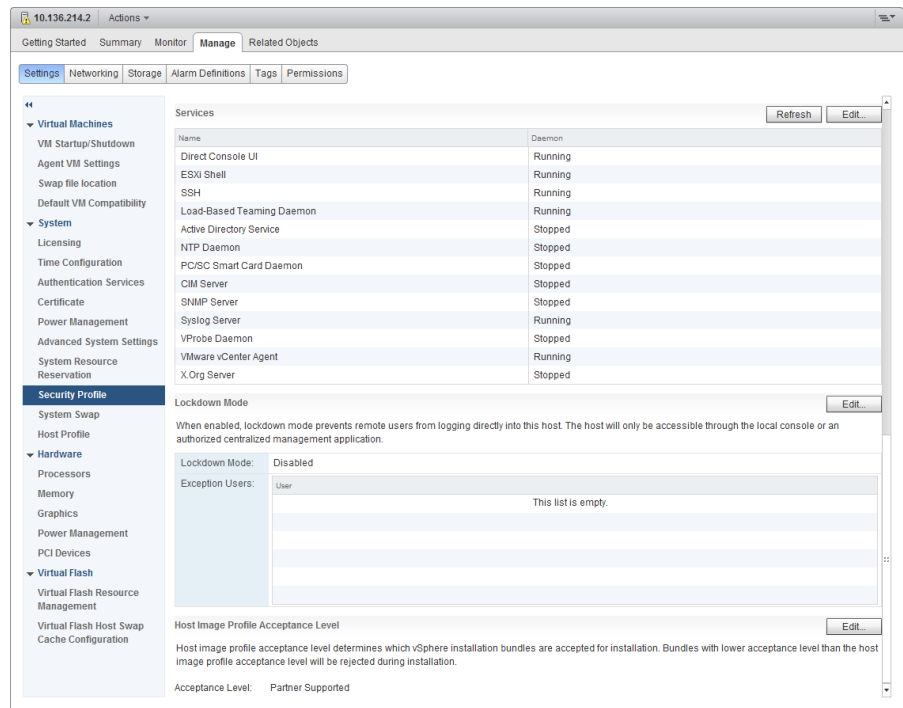
For ease of operation, it is recommended that the parameter value be set as 1 or 2 seconds.

d. If required, update the security parameter values by clicking the **Edit** (✎) icon.

4. Configure the `Lockdown Mode` parameter:

a. In the **vSphere Web Client** screen, under **Manage** menu, select **Settings** > **System** > **Security Profile**.

System-defined parameter panes such as **Services** and **Lockdown Mode**, with predefined parameter options, are displayed in the work area.



- b. Scroll down to access the **Lockdown Mode** pane and verify the value of the **Lockdown Mode** parameter.

Note

The parameter option of **Lockdown Mode** must be **Disabled**. If it is not, change it. To change the parameter option from *Enable* to *Disable*, click **Edit** in the **Lockdown Mode** pane and change the parameter value to *Disable*.

5. Configure the SSH parameter:

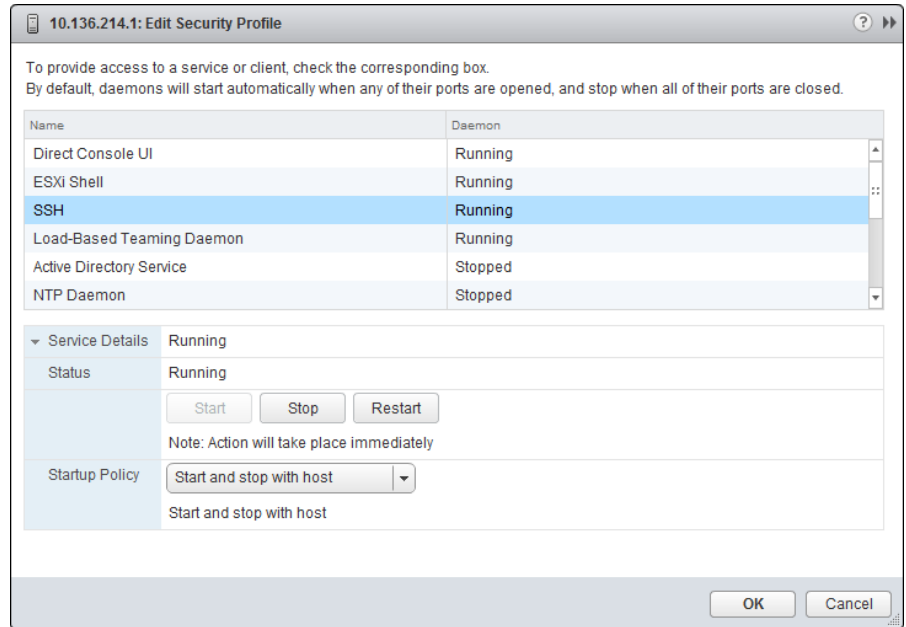
- a. In the **vSphere Web Client** screen, under **Manage** menu, select **Settings > System > Security Profile**.

System-defined parameter panes such as **Services** and **Lockdown Mode**, with predefined parameter options, are displayed in the work area.

- b. Scroll down to access the **Services** pane and verify the SSH option.

Note

SSH must be **Running** (that is, the SSH connection is established). If SSH is **Stopped**, you must change it to **Running**.



c. To change SSH option to Running:

a. Click **Edit** in the **Services** pane.

b. In the **Edit Security Profile** screen, select the SSH parameter.
The SSH status is *Stopped*.

c. Click **Start**.

The SSH status changes to *Running*.

For SSH status as *Stopped*, **Start** is active and **Stop** is grayed out.

d. Click **OK**.

The SSH status changes to *Running*.

6. Restart the management service or restart the server.

CHAPTER 9

Maintaining a ScaleIO System

This chapter describes how to maintain ScaleIO systems. Upgrade procedures are described in the *ScaleIO Upgrade Guide*. Log collection is described in the *ScaleIO Log Collection Technical Notes*. Topics include:

- [Extending the MDM cluster from 3 to 5-node](#).....224
- [Installing RCache on servers in an existing ScaleIO system](#).....227
- [Creating a Lockbox for SNMP, ESRS, or LDAP](#).....228
- [Switching to secured authentication mode](#)..... 228
- [Working with Dynamic Host Name resolution for SNMP in ScaleIO](#).....231
- [Using SCLI in non-secure mode](#)..... 233
- [Extending an existing ScaleIO system](#)..... 233
- [Configuring virtual IP addresses using Installation Manager](#).....235
- [Removing ScaleIO](#).....236

Extending the MDM cluster from 3 to 5-node

This section describes how to extend the MDM cluster from 3-node to 5-node.

Proceed to the section that matches your operating system environment.

Extending the MDM cluster in physical servers

You can extend the MDM cluster in physical servers using the Installation Manager and the CSV topology file.

This topic describes how to deploy ScaleIO on Linux and Windows nodes. Follow the instructions that match your environment and the ScaleIO support matrix for your version.

To extend the MDM cluster, perform the following:

Procedure

1. Get one of the following CSV topology files:
 - The CSV used to deploy the v2.0 system in its current 3-node cluster mode.
 - Download the complete or minimal CSV (provided in the ISO, or can be downloaded from the Installation Manager) and fill-in the current system topology fields in the CSV.
2. Edit, and save the CSV with one of these options:
 - Add two new hosts (two new lines) to the system topology, and in the **Is MDM/TB** column for those lines, designate one as a Slave and one as a Tie Breaker (TB) role.
 - In the **Is MDM/TB** column of two existing hosts that were not part of the MDM cluster, add a Slave and a Tie Breaker (TB) role. This option is displayed in the following figure.

Note

If you need to change the roles of the current nodes, do so only after extending the cluster.

Figure 19 Before extending

	A	B	C	D	E	F	G
1	IPs	Password	Operating System	Is MDM/TB	Is SDS	SDS Device List	Is SDC
2	10.76.60.1	Password1	linux	Primary	Yes	/dev/sdb	Yes
3	10.76.60.2	Password1	linux	Secondary	Yes	/dev/sdb	Yes
4	10.76.60.3	Password1	windows	TB	Yes	g	Yes
5	10.76.60.4	Password1	linux		Yes	/dev/sdb	Yes
6	10.76.60.5	Password1	windows		Yes	g	Yes

Figure 20 After extending

	A	B	C	D	E	F	G
1	IPs	Password	Operating System	Is MDM/TB	Is SDS	SDS Device List	Is SDC
2	10.76.60.1	Password1	linux	Master	Yes	/dev/sdb	Yes
3	10.76.60.2	Password1	linux	Slave	Yes	/dev/sdb	Yes
4	10.76.60.3	Password1	windows	TB	Yes	g	Yes
5	10.76.60.4	Password1	linux	Slave	Yes	/dev/sdb	Yes
6	10.76.60.5	Password1	windows	TB	Yes	g	Yes

3. From the **Packages** tab, upload all ScaleIO packages, per the host OS.
4. From the **Install** tab, select the edited CSV file, and select **Add to existing system** from the drop-down menu.
5. Click **Upload installation CSV**.
6. Start the installation, and monitor as normal.

Converting the CSV topology file after upgrade from ScaleIO v1.32.x

Convert the v1.32.x CSV topology file prior to extending the MDM cluster from 3 nodes to 5 nodes.

Procedure

1. Open the CSV file used to deploy the v1.32.x system.
2. Change the column header from **SDS Pool List** to **StoragePool List**.
3. Replace references to v1.32.x roles with their corresponding v2.x roles.

Instances of "Primary" should be changed to "Master." Instances of "Secondary" should be changed to "Slave."

After you finish

Proceed with [Extending the MDM cluster in physical servers](#) on page 224.

Extending the MDM cluster in VMware servers

This section describes how to extend the MDM cluster in VMware servers. This task is performed by adding two additional ESX servers, and assigning the new roles to them, described in [“Task 1: Extending the MDM cluster”](#).

When you add new manager MDMs, to ensure continued SDC-MDM communication, you should update the SDCs in the system with the new MDM IP addresses. You can do this simply with the vSphere plug-in, described in [“Task 2: Updating the SDC parameters”](#).

Extending the MDM cluster

Procedure

1. If the ESXi servers to be added do not have the SDC component on them, install the SDC on each of the servers, as described in [“Task 3: Installing the SDC on ESX hosts”](#).
2. From the **Basic tasks** section of the screen, click **Deploy ScaleIO environment**.

The ScaleIO VMware deployment wizard begins. If you exited the previous deployment before completion, you will be able to return from where you left off.

NOTICE

The deployment wizard assumes that you are using the provided ScaleIO OVA template to create the ScaleIO virtual machines.

3. In the **Select Installation** screen, select **Add servers to a registered ScaleIO system**, and select the system you want to extend.
4. In the **Select Management Components** screen, perform the following:
5. Select **5-node mode**.
6. In the **Manager MDM** and **Tie Breaker MDM** fields that now appear, select the nodes to add to the cluster.

7. Click **Next**, and continue the deployment.

You can skip steps that do not need to be changed.

Note

When adding components, the wizard adjusts the displayed screens to options that are relevant to the current ScaleIO system.

8. Complete the deployment.

Updating the SDC parameters

This section describes how to use the VMware plug-in to update the SDCs with system parameters that are needed to maintain SDC-MDM communication.

Procedure

1. From the plug-in **Advanced tasks** menu, click **Update SDC parameters**, and follow instructions to complete that process.

2. Ensure that SDC parameters were updated by running this command on each ESX:

```
cat /etc/vmware/esx.conf |grep scini|grep -i mdm
```

Installing RFCache on servers in an existing ScaleIO system

This section describes how to install RFCache on servers in an existing ScaleIO system.

Proceed to the section that matches your operating system environment, physical or virtual.

Installing RFCache in physical servers

Installing RFCache in physical servers is performed with the Installation Manager and the CSV topology file, as described in [“Extending an existing ScaleIO system”](#).

In a Windows installation, the servers need to be restarted after installing.

To install RFCache, perform the following:

Procedure

1. Update the CSV file with the RFCache fields, as described in the "Preparing the CSV topology file" section of the *ScaleIO Deployment Guide*.
2. Follow the instructions in [“Extending an existing ScaleIO system”](#).

Enabling RFCache on VMware servers

Enabling RFCache on VMware servers is performed manually and with the vSphere plug-in. First, you copy the package to the SVMs, and then use the plug-in to configure its use.

Procedure

1. Copy the RFCache package (`xcache`) to the SVMs.
2. Enable RFCache on every SDS server that is to provide acceleration:
 - a. Copy the following file to all v2.0 SVMs running an SDS:

```
EMC-ScaleIO-xcache-x.x-x.0.slesxx.x.x86_64
```

- b. Install the file:

```
rpm -i EMC-ScaleIO-xcache-x.x-x.0.slesxx.x.x86_64
```

3. Use the vSphere plug-in to enable RFCache on the SDS:
 - a. Click **SDSs**.
 - b. Right-click an SDS, and select **Add devices to a single SDS**.
 - c. Click a device in the **Use for** drop-down, select **RFcache**.

d. Click **OK**.

4. Repeat the previous step for every SDS on which you want to enable RFcache.

Creating a Lockbox for SNMP, ESRS, or LDAP

A Lockbox file must exist on the ScaleIO Gateway for use by the SNMP, ESRS, and LDAPS (secure LDAP) features. A lockbox file is optional for LDAP usage. The Lockbox is used to securely store MDM authentication credentials. SioGWTool is used to create the Lockbox and to add MDM authentication credentials.

For more information about configuring a Lockbox, and about other SNMP configuration activities, see “Configuring SNMP properties after deployment” in the *ScaleIO User Guide*.

For more information about configuring a Lockbox for ESRS, and other ESRS configuration activities, see "Configuring ESRS connection properties" in the *ScaleIO Deployment Guide*.

For more information about LDAP, see the *User Roles and LDAP Usage Technical Notes*.

Switching to secured authentication mode

This section describes how to switch the authentication mode of a ScaleIO system from non-secured to secured mode.

This can be relevant when upgrading to 2.0, where the upgraded system is always in non-secure mode, or after performing a clean install with the Installation Manager and choosing non-secure authentication, via the advanced installation options.

Secure authentication can only be enabled on operating systems that support Open SSL v1.0.1 or later.

You need to update the following areas:

- Internal secure mode (between the MDM and SDS)
- External secure mode (between the management clients and MDM)
- LIA authentication (between the Installation Manager and the LIA)

Proceed to the section that describes your environment:

- [“Physical Linux servers”](#)
- [“Physical Windows servers”](#)
- [“VMware servers”](#)

Physical Linux servers

This topic describes how to change the authentication mode from non-secure to secure on physical Linux servers.

Before beginning, verify the version of the OpenSSL on the server, by running this command:

```
rpm -qa | grep -i openssl
```

If the version is 1.0.1 or greater, you can continue. If it is not, secure authentication cannot be enabled.

Procedure

1. Log in to the SCLI, in a non-secure fashion:

```
scli --login --username <USER> --password <PASSWORD> --
use_nonsecure_communication
```

2. Change to internal secure mode, by running the following command:

```
scli --set_component_authentication_properties --
use_authentication --use_nonsecure_communication
```

3. Change to external secure mode, by running the following command:

```
scli --set_management_client_communication --
enable_client_secure_communication --
use_nonsecure_communication
```

4. Enable secure LIA authentication, by performing the following:

Note

This procedure can take quite some time, and there is no roll-back.

- a. Log in to the IM server (https://<IM_Server_IP>).
- b. Click the **Maintain** tab.
- c. From the **Maintenance operation** screen, perform the following:
 - a. Type the Master (or Primary) MDM credentials.
 - b. Type the LIA password.
 - c. In **Advanced Options**, select **Allow non-secure communication with LIAs**.
 - d. Click **Retrieve system topology**.
 - e. If a certificate approval is required, approve it.
 - f. Type the authentication credentials again.
 - g. Click **Retrieve system topology**.
 - h. Click **Security settings** and select **Enable LIAs security**.
 - i. Type the authentication credentials again.
 - j. Confirm the request, and click **Enable LIAs security**.
 - k. If a certificate approval is required, approve it.
 - l. When complete, click **Mark operation completed**.

Physical Windows servers

This topic describes how to change the authentication mode from non-secure to secure on physical Windows servers.

Before beginning, verify (in **Programs and Features**) that OpenSSL version 1.0.1 or greater is installed. If it is not, secure authentication cannot be enabled.

Procedure

1. Log in to the SCLI, in a non-secure fashion:

```
scli --login --username <USER> --password <PASSWORD> --
use_nonsecure_communication
```

2. Change to internal secure mode, by running the following command:

```
scli --set_component_authentication_properties --
use_authentication --use_nonsecure_communication
```

3. Change to external secure mode, by running the following command:

```
scli --set_management_client_communication --
enable_client_secure_communication --
use_nonsecure_communication
```

4. Enable secure LIA authentication, by performing the following:

Note

This procedure can take quite some time, and there is no roll-back.

- a. Log in to the IM server (https://<IM_Server_IP>).
- b. Click the **Maintain** tab.
- c. From the **Maintenance operation** screen, perform the following:
 - a. Type the Master (or Primary) MDM credentials.
 - b. Type the LIA password.
 - c. In **Advanced Options**, select **Allow non-secure communication with LIAs**.
 - d. Click **Retrieve system topology**.
 - e. If a certificate approval is required, approve it.
 - f. Type the authentication credentials again.
 - g. Click **Retrieve cluster topology**.
 - h. Click **Security settings** and select **Enable LIAs security**.
 - i. Type the authentication credentials again.
 - j. Confirm the request, and click **Enable LIAs security**.
 - k. If a certificate approval is required, approve it.
 - l. When complete, click `Mark operation completed`.

VMware servers

This topic describes how to change the authentication mode from non-secure to secure on VMware servers.

To enable secure authentication mode, perform the following:

Procedure

1. Log in to the SCLI, in a non-secure fashion:

```
scli --login --username <USER> --password <PASSWORD> --
use_nonsecure_communication
```

2. Change the internal secure mode, by running the following command:

```
scli --set_component_authentication_properties --
use_authentication --use_nonsecure_communication
```

3. Change the external secure mode, by running the following command:

```
scli --set_management_client_communication --
enable_client_secure_communication --
use_nonsecure_communication
```

4. Enable secure LIA authentication, by performing the following:

- a. Open the ScaleIO Gateway Installation Manager:

- a. From the plug-in **ScaleIO Systems** screen, right-click the system for which you want to perform this operation.

- b. Select **Open ScaleIO Gateway**.

Note

Opening the Gateway can take some time.

- c. From the **IM welcome** screen, enter the default IM credentials.

- b. From the **IM web client** main menu, click **Maintain**.

- c. In **Advanced Options**, select **Allow non-secure communication with LIAs**.

The **Maintenance operation** screen appears.

- a. Click **Security settings** and select **Enable LIAs security**.

- b. Confirm the request, and click **Enable LIAs security**.

- c. If a certificate approval is required, approve the certificates.

5. From the main plug-in window, click **ScaleIO systems**.

6. Right-click the system to register, and select **Reregister ScaleIO system**.

7. Enter the user name and password of the system, and click **OK**.

8. Approve the Master MDM certificate by clicking **Accept**.

9. From the plug-in **ScaleIO Systems** screen, right-click the system and accept the Slave MDM certificate.

Note

Before performing maintenance operations in a system where replication is enabled on SDC nodes, you should exclude the RPA nodes from the detection list, as described in [“Configuring the Installation Manager”](#).

Working with Dynamic Host Name resolution for SNMP in ScaleIO

When working with Dynamic Host Name resolution and SNMP, specific configurations are required in the ScaleIO Gateway and on the DNS server.

Before you begin

Ensure that SNMP is already enabled and that SNMP traps are being received by at least one SNMP trap receiver.

The following procedures are required in order to set up Dynamic Host Name resolution:

- On the ScaleIO Gateway, add the host name of the SNMP trap receiver to the appropriate parameter in the `gatewayUser.properties` file
- On the DNS server, configure the SNMP trap receiver properties, in order to support dynamic host name resolution
- On the DNS server, reduce the "Time To Live" (TTL) setting for the SNMP trap receiver

Procedure

1. On the ScaleIO Gateway:
 - a. Open the `gatewayUser.properties` file in a text editor. The file location is:
 On Linux: `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/gatewayUser.properties`
 On Windows: `C:\Program Files\EMC\ScaleIO\Gateway\webapps\ROOT\WEB-INF\classes\gatewayUser.properties`
 - b. Add the host name of the SNMP trap receiver to the property `snmp.traps_receiver_ip=`. If there is more than one IP address or host name, use a comma-separated list.
 - c. Save the file.
 - d. Restart the `scaleio-gateway` service.
 - e. Verify that traps are being received at all configured trap receiver hosts.
2. On the DNS server, configure dynamic host name resolution support for the trap receiver.
3. On the DNS server, reduce the TTL setting for the trap receiver. For example, on Windows, perform the following:
 - a. Open the **DNS manager** window.
 - b. Click **View**, and select the **Advanced** option.
 - c. Right-click the trap receiver, and select **Properties**. A window similar to the following is displayed:

d. At the bottom of the window, in the **Time to live (TTL)** field, change the value from one hour to one or two seconds.

e. Click **OK**.

Results

Dynamic Host Name resolution configuration is complete.

Using SCLI in non-secure mode

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

Procedure

1. On every MDM server, disable secure communication:
 - Windows - In the SCLI `conf.txt` file, add `cli_use_secure_communication=0`
 - Linux - Run `echo cli_use_secure_communication=0 >> ~/.scli/conf.txt`

Extending an existing ScaleIO system

This section describes how to add components to an existing ScaleIO installation.

In physical environments, you add components with the Installation Manager. In VMware environments, you add components with the VMware deployment wizard.

Adding components with the Installation Manager

This section describes how to add components with the Installation Manager.

To add components, you first need to update the CSV topology file with the new components, then you can use the web client to add them.

The secure communication mode of components must match the mode of the system to which the components are being added. If they do not match, you must either change the mode of the components to be added or change the mode of the system so that they match.

Procedure

1. Follow the procedure described in [“Installing with the Installation Manager”](#).

Note

Use the same LIA password that was configured during initial installation.

2. In the Upload CSV stage, browse to the updated CSV file, and select **Add to existing sys.**
3. Upload the CSV, and continue as normal.

Adding components with the VMware deployment wizard

This section describes how to add ScaleIO components to an existing system with the VMware deployment wizard.

Note

The following procedure cannot be used to add an SDS component to an existing SVM. To do so, contact EMC Support.

Procedure

1. From the **Basic tasks** section of the screen, click **Deploy ScaleIO environment**.

The ScaleIO VMware deployment wizard begins. If you exited the previous deployment before completion, you will be able to return from where you left off.

NOTICE

The deployment wizard assumes that you are using the provided ScaleIO OVA template to create the ScaleIO virtual machines.

2. In the **Select Installation** screen, select **Add servers to a registered ScaleIO system**, and select the system you want to extend.
3. Continue with the deployment steps, adding the new nodes.
You can skip steps that do not need to be changed.

Note

When adding components, the wizard adjusts the displayed screens to options that are relevant to the current ScaleIO system.

4. Complete the deployment.

Note

After extending an existing SVM with a new ScaleIO role/component, you must perform manual memory allocation on the SVM, as described in [“SVM manual memory allocation”](#).

Configuring virtual IP addresses using Installation Manager

Configure virtual IP addresses using the **Maintain** menu in the Installation Manager (IM).

You can assign a virtual IP address for each possible manager MDM, which will be used for communications between the MDM cluster and SDCs. Only one virtual IP address can be mapped to each NIC, with a maximum of four virtual IP addresses per system. The IM can be used to assign new virtual IP addresses only; to change or remove existing virtual IP addresses, use the appropriate CLI commands.

Virtual IP addresses are not supported on nodes using Windows operating system.

Procedure

1. In the Installation Manager, select **Maintain**.
2. Select **Set Virtual IPs**.
3. In the **Set Virtual IPs for ScaleIO system** screen, type the MDM password.

Figure 21 Set Virtual IPs for ScaleIO system screen

Virtual IPs:	MDM Virtual IP Interfaces:
192.168.111.7	eth4
192.168.111.8	eth5
10.103.110.7	eth4
10.103.110.172	eth5

4. For each MDM that you wish to set a virtual IP address, enter a virtual IP address and the NIC to which it will be mapped. For each new virtual IP address,

enter the virtual IP address and NIC name for each MDM to which it will be mapped.

With the IM, you can configure NIC names that contain the following characters only: a-z, A-Z, 0-9. If a NIC name contains the "-" or "_" character (for example eth-01), don't use the IM. Configure this IP address with the CLI

`modify_virtual_ip_interfaces` command and the `--new_mdm_virtual_ip_interface <INTF>` parameter.

5. Click **Set Virtual IPs**.

Results

The virtual IP address is configured and all of the SDCs are updated with the new virtual IP address.

Removing ScaleIO

This section describes how to remove ScaleIO.

To uninstall ScaleIO, use the Installation Manager. This requires that the LIA be installed in all nodes to be changed.

When removing RFcache (the `xcache` package) on a Windows server, a server restart is necessary after the removal.

To unregister the vSphere plug-in, see "Unregistering the ScaleIO plug-in".

Removing ScaleIO using the IM

This section describes how to use the Installation Manager to remove ScaleIO. All ScaleIO components in the system that is being accessed will be removed. This information is attained from the LIA that is installed on every node.

Procedure

1. Log in to the web client, as described in ["Installing with the Installation Manager"](#).
2. From the IM web client main menu, click **Maintain**.
3. In the **Maintenance operation** screen, type the authentication credentials, then click **Retrieve system topology**.

The system topology is displayed.

4. Click the **Show Uninstall button** link, and confirm enabling this option.

The uninstall operation may take some time, depending on your system topology. This operation cannot be rolled back.

5. Click **Uninstall**.

A confirmation dialog is displayed.

NOTICE

Uninstalling an SDC component requires a machine restart. If you are uninstalling SDC components on Windows servers, select to enable automatic restart (on those servers only). Alternatively, you can manually restart these servers after removing the SDC.

On Linux servers, if the kernel module is busy, perform a manual restart.

6. Enter the MDM password, select to reboot servers (optional), and click **Uninstall**.
7. To monitor the uninstallation progress, click **Monitor**.
8. When the uninstallation is complete, click **Mark operation completed**.

Unregistering the ScaleIO plug-in

To remove the currently register plug-in, perform the following:

Procedure

1. Run the script to remove the plug-in:
 - a. From the folder where you extracted the current vSphere web plug-in ZIP file (for example: `EMC-ScaleIO-vSphere-web-plugin-package-2.0-XXX.X.zip`), use PowerCLI to run the ScaleIO plug-in script (for example: `ScaleIOPluginSetup-2.0-XXX.X.ps1`).
 - b. Select option 2, `Unregister ScaleIO plugin`.
2. Enter the vCenter credentials and confirm the script actions.
3. Log out, then log back in to the vSphere web client.
4. The plug-in is no longer registered.

CHAPTER 10

System Analysis

This chapter describes the ScaleIO system analysis tool. Topics include:

- [System analysis overview](#)240

System analysis overview

This topic describes ScaleIO system analysis and the environment required to use it.

ScaleIO system analysis enables you to identify potential issues with your ScaleIO system which may prevent best performance. It is highly recommended to analyze the ScaleIO system immediately after deployment, before provisioning volumes, and before using the system in production. You can also use it to check the health of a system that is already operational.

The system analysis is invoked from the ScaleIO Installation Manager (IM). The analysis checks the following:

- ScaleIO components are up and running
- Ping between two relevant nodes in the system
- Connectivity within the ScaleIO configuration (for example, connectivity between SDSs within a Protection Domain, connection of SDCs with the cluster virtual IP address)
- Network configuration
- RAID controller and device configuration

Using the system analysis, you can detect any potential issues in the system, then rectify them, before provisioning and using the system in a production environment.

Environment requirements and prerequisites:

- Supports servers running on the following operating systems:
 - RHEL 6.x, 7.x
 - SLES 11.3, 12.0, 12.1, 12.2
 - ESX 5.5, 6.0, 6.5
- All servers must have the following 3rd-party tools installed on them:
 - Netcat
 - StorCLI or PercCLI
 - smartctl
- Requires a ScaleIO Gateway server:
 - On a Linux or Windows server; ScaleIO is tested on RHEL 6.x and 7.x, on SLES 11.3, 12.0, 12.1, and 12.2, and on Windows Server 2012.
 - At least 1 GB of free disk space per node in the system.
- A web browser, that is supported by the IM.
- Supports LSI RAID controller cards.
- Supports IPv4/IPv6 network configuration.
- Supports iproute2. Newer operating systems, such as RHEL 7.x, use iproute2 commands.

Best-practice recommendation:

1. Deploy ScaleIO.
2. Analyze system to identify issues that should be fixed.
3. Fix issues.

4. Analyze system to verify that the issues have been fixed.
5. When the system meets your satisfaction, you can move it into production.

Limitations and compatibility:

- On servers with MegaRAID, the RAID function analysis is not supported; only StorCLI is supported.
- On servers with multiple SDSs, analysis is performed but it is not conclusive.
- RFCache analysis is not performed.
- Due to default Internet Explorer settings, to expand a report, you may need to grant permission for IE to run scripts. For more information, see the version release notes.

Creating the system analysis report

The topic describes how to create and display the ScaleIO system analysis report.

Before you begin

Ensure that you have access to the following:

- A web browser that is supported by the ScaleIO Installation Manager (IM)
- IP address of the ScaleIO Gateway server
- The Gateway admin username (default: admin) and password (defined during deployment)
- Master MDM IP address, username, and password (defined during deployment)
- LIA password

Note

In 2-layer systems where SDCs are ESX-based, no LIA passwords are required.

- Root password (for ESX-based servers)
- IP addresses of ESX-based SDCs in your network that are running ScaleIO versions earlier than version 2.0.1

Procedure

1. Log in to the IM server:
 - a. Point your browser to **https://<IM_Server_IP_address>**
 where **<IM_Server_IP_address>** is the IP address of the server on which the Gateway package is installed.
 The **ScaleIO Installer** login screen is displayed.
 - b. If a login banner is displayed, accept it to continue.
 - c. Type the Gateway username (default is admin) and password.
 - d. Click **Login**.
 The **ScaleIO Installer** screen is displayed.
2. Generate the system analysis report:
 - a. Click the **Maintain** tab.
 The **Maintenance operation** screen is displayed.
 If a warning is displayed that indicates that a previous operation is not completed, it means that the last IM operation was not marked as complete.

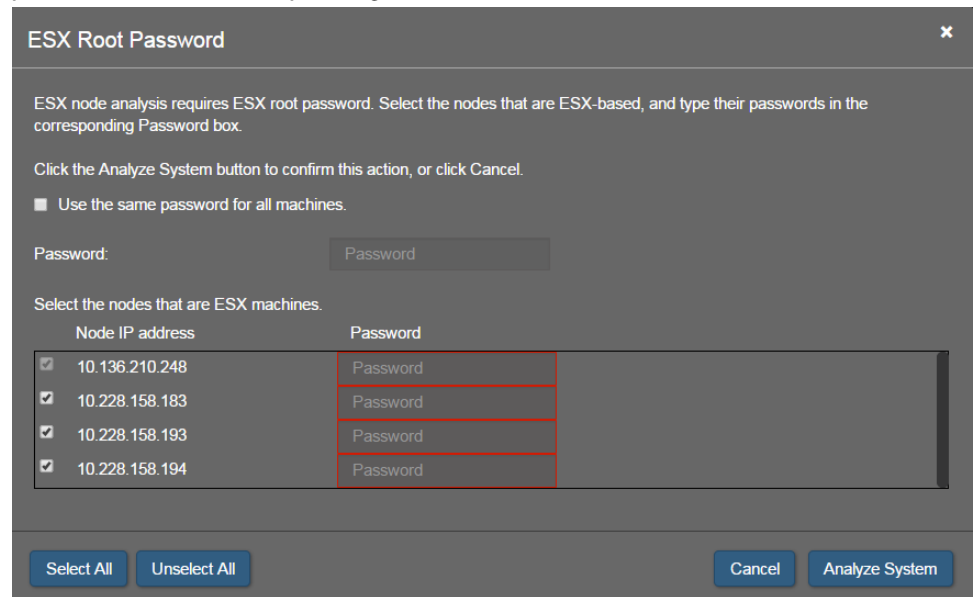
To continue, click the **Monitor** tab, mark the operation complete, then return to the **Maintain** tab.

b. Enter the authentication credentials:

- Master MDM IP address, user name, and password
- LIA password (Linux-based systems only)

c. Click **Retrieve system topology**.

If ESX-based servers exist in the system, the following window appears. ESX servers that are automatically recognized are already selected. Other servers, on which SDCs from earlier ScaleIO versions are installed, may also be shown. Select the ones that are ESX-based, and type their root passwords in the corresponding boxes. Click **Analyze System**.



ESX Root Password

ESX node analysis requires ESX root password. Select the nodes that are ESX-based, and type their passwords in the corresponding Password box.

Click the Analyze System button to confirm this action, or click Cancel.

☐ Use the same password for all machines.

Password:

Select the nodes that are ESX machines.

Node IP address	Password
<input checked="" type="checkbox"/> 10.136.210.248	<input type="password"/>
<input checked="" type="checkbox"/> 10.228.158.183	<input type="password"/>
<input checked="" type="checkbox"/> 10.228.158.193	<input type="password"/>
<input checked="" type="checkbox"/> 10.228.158.194	<input type="password"/>

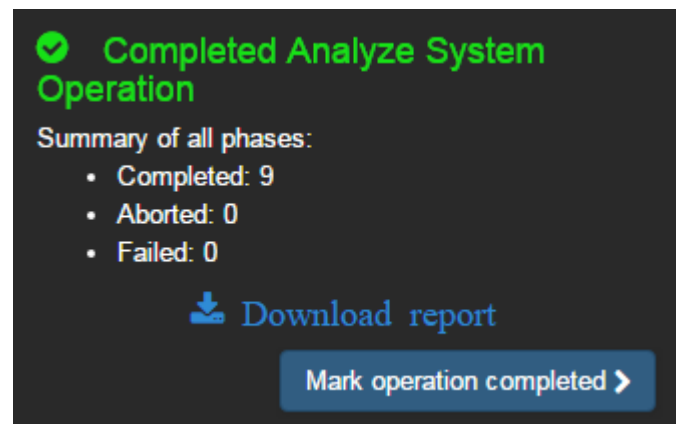
The system topology is displayed.

d. Click **System Log & Analysis** then **Analyze System**.

e. Monitor the progress of the system analysis by clicking the **Monitor** tab.

When the process completes (which could take several minutes), a **Download report** link is displayed.

f. Click the **Download report** link.



Completed Analyze System Operation

Summary of all phases:

- Completed: 9
- Aborted: 0
- Failed: 0

[Download report](#)

The report is saved, in ZIP format, in the default download location of the server on which the report was run.

g. To enable the IM to be available for subsequent operations, click **Mark operation completed**.

3. Display the analysis report:

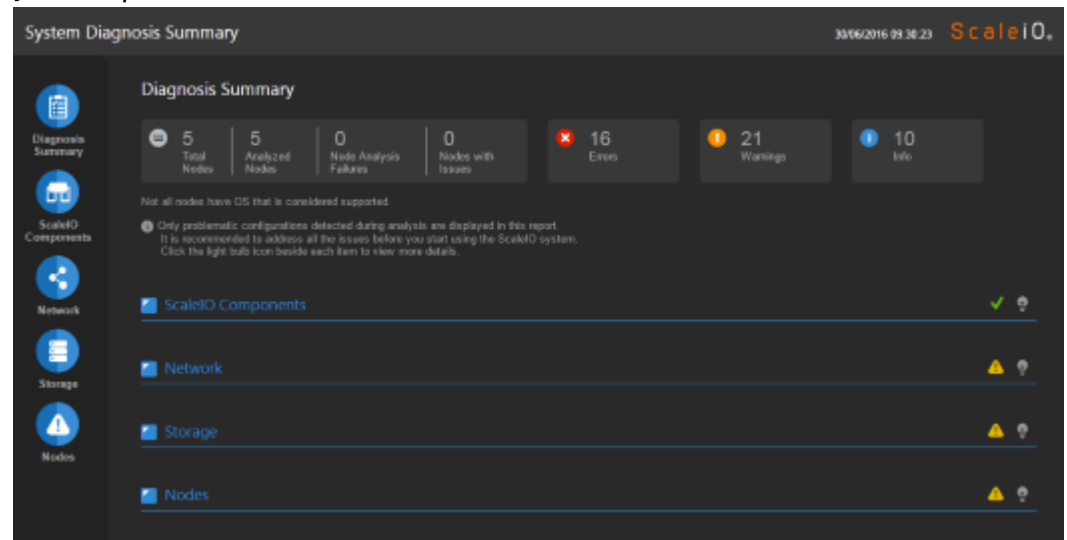
a. Unzip the downloaded file.

The file includes the `ScaleioSystemDiagnosisReport.html` analysis file, and several TGZ files (one for each node, in the `dumps` folder).

b. Double-click `ScaleioSystemDiagnosisReport.html`.

Results

The **System Diagnosis Summary** report is displayed in the default web browser on your computer.

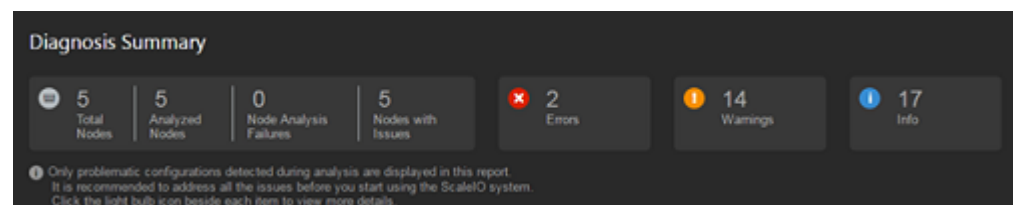


System analysis report description

This topic describes how to get the most benefit from the system analysis report.

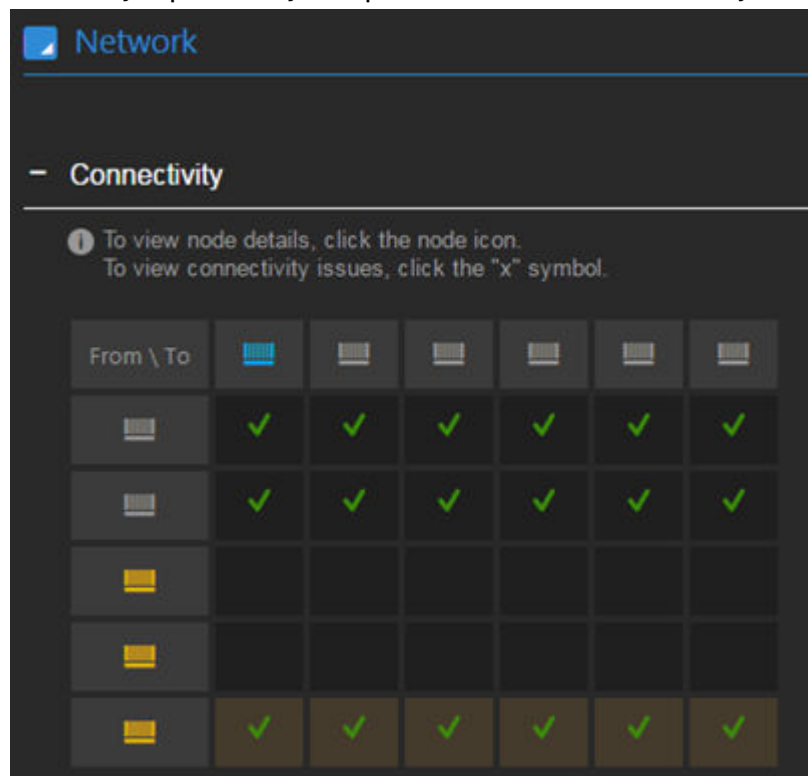
At the top of the report, the Diagnosis Summary shows the number of the issues that have been detected system-wide. The summary shows the following categories:

- Node analysis
 - How many nodes are analyzed
 - How many nodes could not be analyzed
 - How many nodes have issues
- Severity analysis
 - How many issues of each severity level were found



When the analysis first opens, the major sections are shown in summary form. You can expand them to show detailed diagnostic reports, as follows:

- **ScaleIO components:**
This section of the report shows the non-running ScaleIO server components, that is, SDS, SDC, and MDM. Failure of these components may affect system performance and data availability. Each of the ScaleIO server components supports the following functionality:
 - **SDS server**
The SDS (ScaleIO Data Server) manages the capacity of a single server and acts as the back-end for data access. The SDS is installed on all the servers contributing storage devices to the ScaleIO system. Failure of an SDS may affect the cluster performance and data availability.
 - **SDC server**
The SDC (ScaleIO Data Client) server is installed on each server that needs access the ScaleIO storage and it is the gateway to the ScaleIO storage. Failure of an SDC server denies its application access to the ScaleIO storage.
 - **MDM server**
The MDM (Meta Data Manager) server controls and monitors the ScaleIO system. Failure of an MDM may affect the cluster performance and data availability.
- **Network:**
This section of the report checks the connectivity between various ScaleIO components, as well as the NIC configuration and performance. The network issues may impact the system performance and data availability.



- **Connectivity**
Performing pings between ScaleIO components leads to detecting and resolving connectivity-related issues in the system. If the regular pings succeed, then the MTU pings, followed by the Netcat pings are performed to the ports used by the ScaleIO application.

If virtual IP addresses are assigned to the MDMs in the cluster, a logical host, called *MDM cluster* is displayed (represented by a blue host icon) in the analysis report. The following issues are tested:

- SDC connectivity with the virtual IP addresses.
- All virtual IP addresses are configured.
- Only one physical node is configured to use the virtual IP address.

To view specific information:

- To view the node details, click the node icon.
- To view connectivity issues, click the X symbol

- **NIC Configuration and Performance**

This section of the report displays the configurations that do not meet the best practice recommendations described in the user documentation.

- **MDM cluster**

Shows status of virtual IP addresses assigned to the MDMs in the cluster.

- **Storage:**

The Storage section of the report describes the issues associated with the RAID controllers, storage devices, and Storage Pool configurations.

- **RAID Controller**

Describes the issues related to the physical disks

- **Devices**



Describes the list of problematic or potentially problematic storage devices

- **Storage Pool Uniformity**

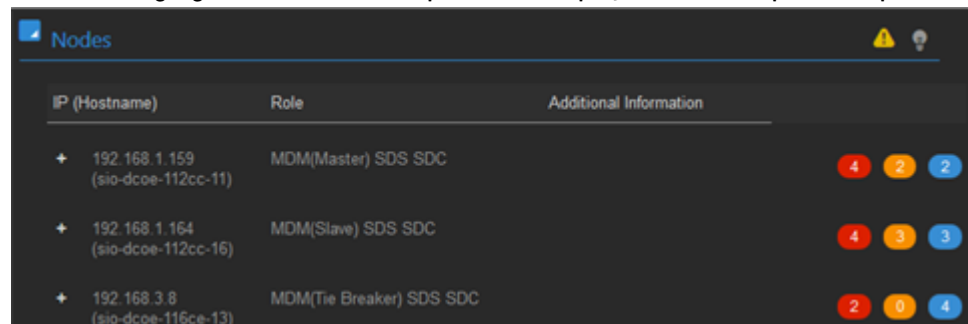
Indicates Storage Pools with non-homogeneous disk performance

- **Nodes (groups all of the reported issues per node):** Describes the detected issues, as listed above, grouped by node


You can show more (or less) information, as follows:


- Use the Expand () button and + symbol to drill down.
- Use the Collapse () button and – symbol to hide information.

The following figure shows an example of the display when the report is expanded.



IP (Hostname)	Role	Additional Information
+ 192.168.1.159 (sio-dcoe-112cc-11)	MDM(Master) SDS SDC	4 2 2
+ 192.168.1.164 (sio-dcoe-112cc-16)	MDM(Slave) SDS SDC	4 3 3
+ 192.168.3.8 (sio-dcoe-116ce-13)	MDM(Tie Breaker) SDS SDC	2 0 4

You can show additional details of an error by clicking the () icon. A pop-up window similar to the following is displayed.











Diagnosis
NICs with auto-negotiation off.

Business Impact
The wrong setting might cause severe performance issues: lower IOPS and throughput.

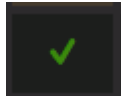
Suggested Solution
Enable NIC auto-negotiation.



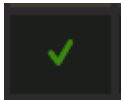
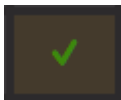




Severity
Warning

The following table describes how to navigate the analysis information:

To display this...	Click this symbol...
Diagnosis Summary	
ScaleIO component issues, sorted according to SDS, SDC, and MDM.	
Network connectivity (including virtual IP address configuration) and NICs	
Storage, including RAID controller, devices, and Storage Pool uniformity	
Nodes, displayed according to IP address and system role	
Collapse a list	
Expand a list	
Open a pop-up containing Diagnosis, Business Impact, and Suggested Solution information.	

The following table describes other symbols and interface elements used in the analysis report display:

Symbol / Interface Elements	Description
	All connectivity tests passed

Symbol / Interface Elements	Description
 or 	Connectivity issues were detected (for more information, click the symbol)
 (black background)	Connectivity matrix background: all pertinent connectivity tools are available
 (brown background)	Connectivity matrix background: connectivity tests could not be performed, due to a missing tool
	Error counters: <ul style="list-style-type: none"> • Red - Error • Orange - Warning • Blue - Info
	All connectivity test tools are available for use on the node
	Some connectivity tests could not be performed on this node, due to missing tools (for more information, click the symbol)
	A logical host, that represents the MDM cluster, when virtual IP addresses are configured

Note

An empty cell in the connectivity matrix indicates that no connectivity check was performed. In such cases, no connectivity is expected.

CHAPTER 11

Configuring ESRS connection properties

This chapter explains how to configure connection properties, to allow communication between the ScaleIO system and ESRS (EMC Secure Remote Support) servers. ESRS can be used by ASP Partners for remote hardware support. For more information about using ESRS, refer to ESRS documentation.

Use SioGWTool in command line to configure the connection properties.

Topics in this chapter include:

- [Before configuring ESRS](#)..... 250
- [Registering the system with the ESRS Gateway](#).....250
- [Performing other ESRS configuration activities](#).....252
- [Validating ConnectEMC Dial-Home](#).....253

Before configuring ESRS

Before configuring ESRS, ensure that your system meets the following requirements, and that you have the following information:

- Requirements:
 - ESRS Gateway v3 version 3.08 or higher must be installed and configured.
 - ESRS Gateway must be reachable from the ScaleIO-Gateway node on port 9443.
 - The ScaleIO license must be installed.
 - The ScaleIO Master MDM Management IP address to be used as the Connect-In IP address must be an IP address that is accessible from the ESRS Gateway (for example, in case of NAT). Ensure that you know this IP address, because you will need it when you perform the registration procedure.
 - On the ScaleIO Gateway, in the `gatewayUser.properties` file, set the property `features.enable_esrs=true`
- Information:
 - One or more IP addresses of the ESRS servers.
 - ESRS user and password credentials for the ScaleIO command.
 - MDM user name and password.

Registering the system with the ESRS Gateway

Enable use of the ESRS feature on your ScaleIO system, add the ESRS Gateway's certificate to the truststore, and register your system on the ESRS Gateway.

Before you begin

See [Before configuring ESRS](#) on page 250.

The REST procedures described below are performed using cURL. You may use other similar tools and their corresponding commands to perform those steps.

A web browser is required for this procedure.

In order to work with ESRS, the LockBox must be configured, and the MDM credentials must be added to it. The LockBox is required for both SNMP and ESRS. The same LockBox is used for both functionalities. Use SioGWTool in command line to create a LockBox, to register the ESRS gateway address, and to set the connect-in IP address. SioGWTool is located in:

- Linux:


```
/opt/emc/scaleio/gateway/bin/SioGWTool.sh
```
- Windows:


```
C:\Program Files\EMC\ScaleIO\Gateway\bin\SioGWTool.bat
```

Procedure

1. If no LockBox exists on the ScaleIO system (one may already have been created for SNMP, for example), create a LockBox. In command line, use the following SioGWTool commands:

```
<SioGWTool_PATH> --change_lb_passphrase --new_passphrase  
<new_passphrase>
```

```
SioGWTool --set_mdm_credentials --mdm_user <mdm_user_name>
--mdm_password <mdm_password>
```

2. Enable the ESRS feature in the `gatewayUser.properties` file.

- a. Use a text editor to open the `gatewayUser.properties` file, located in the following directory on the Installation Manager/ScaleIO Gateway server:

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

- b. Change the `enable_esrs` feature to "true":

```
features.enable_esrs=true
```

- c. Save the file.

3. Add the ESRS Gateway's certificate to the truststore.

- a. In a web browser, browse to `<ESRS_Gateway_IP_address>:9443`.
- b. Download the certificate that is displayed, and save it as a file.
- c. Log in to REST and get a token. Make a note of the token.

```
curl -k -v --basic --user
admin:<mdm_admin_password> https://
<ScaleIO_Gateway_IP_address>/api/login
```

- d. Using REST, add the certificate to truststore.

```
curl -k -v --basic -
uadmin:<token_received_from_previous_command> --form
"file=@<path_to_certificate_file>" https://
<ScaleIO_Gateway_IP_address>/api/trustHostCertificate/Mdm
```

4. Restart the `scaleio-gateway` service.

Option	Description
Windows	Restart the EMC ScaleIO Gateway service
Linux	Type the command <pre>service scaleio-gateway restart</pre>

5. Register the ScaleIO system on the ESRS Gateway. In command line, use the following `SioGWTool` command:

```
<SioGWTool_PATH> --register_esrs_gateway --scaleio_gateway_ip
<IP_ADDRESS:PORT> --scaleio_gateway_user <USER> --
```

```
scaleio_gateway_password <PASSWORD> --esrs_gateway_ip
<IP_ADDRESS> --esrs_gateway_user <USER>
--esrs_gateway_password <PASSWORD> --connect_in_ip <IP_ADDRESS>
```

where `--esrs_gateway_user` is the user name used for EMC support (typically, an email address), and `--esrs_gateway_password` is its corresponding password. `--connect_in_ip` is the MGMT IP address of the current Master MDM.

Results

ESRS registration is complete.

Performing other ESRS configuration activities

Use SioGWTool in command line for ESRS configuration activities.

SioGWTool is located in:

- **Linux:** /opt/emc/scaleio/gateway/bin/SioGWTool.sh
- **Windows:** C:\Program Files\EMC\ScaleIO\Gateway\bin\SioGWTool.bat

Procedure

1. The following activities can be performed with SioGWTool:

Activity	Command
To unregister from all ESRS Gateways, type the command:	<SioGWTool_PATH> --unregister_esrs_gateway --scaleio_gateway_ip <ip address:port> --scaleio_gateway_user <user> --scaleio_gateway_password <password>
To remove a specific ESRS Gateway from the ScaleIO Gateway configuration, type the command:	<SioGWTool_PATH> --remove_esrs_gateway --scaleio_gateway_ip <scaleio_gateway_ip_address> --scaleio_gateway_user <scaleio_gateway_user> --scaleio_gateway_password <scaleio_gateway_password> --esrs_gateway_ip <esrs_gateway_ip_address>
To reset the ScaleIO Gateway admin password, type the command:	<SioGWTool_PATH> --reset_password
To start ESRS logic on the ScaleIO Gateway (alerts will be sent to the ESRS server), type the command:	<SioGWTool_PATH> --start_esrs
To stop ESRS logic from running on the ScaleIO Gateway (alerts will not be sent to the ESRS server), type the command:	<SioGWTool_PATH> --stop_esrs

Activity	Command
To check the connection to a registered ESRS Gateway (probing will be done throughout the servers to find one that can be connected), type the command:	<SioGWTool_PATH> --check_esrs_connectivity
To show a list of the available SioGWTool commands, type the command:	<SioGWTool_PATH> --help

Validating ConnectEMC Dial-Home

Verify that ConnectEMC (ESRS) is correctly configured and is communicating with Dell EMC Global Services, by performing a dial-home test.

This procedure describes how to quickly verify the connectivity between the ESRS-GW and ESRS servers, either after installation or after registering to ESRS. The ESRS-GW is located at the customer's site and has a verified connection to ESRS servers as part of its installation.

Procedure

1. While logged in to the ESRS-gateway as the gpadmin operating system user, remove the contents of the directory `/opt/connectemc/archive`.
2. Create a test event for ConnectEMC using the following command:

```
touch /opt/connectemc/poll/testerror
```

Results

Within up to 10 minutes (depending on your polling interval setting), ConnectEMC (ESRS) will generate a file in one of the following directories, depending on the test result:

If the connection was successful, a file will be saved in `/opt/connectemc/archive`.

If the service fails, a file will be saved in `/opt/connectemc/archive/failed`.

CHAPTER 12

Common Tasks

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

• Install the ScaleIO GUI	256
• Log in to the ScaleIO GUI	256
• Connection and disconnection information	257
• Add LIA to a system to enable automated upgrade	257
• Associating ScaleIO volumes with physical disks	258
• Port usage and changing default ports	261
• Adding an external SDC to an existing ScaleIO system	262
• Changing the LIA configuration file	265
• Cleaning the ScaleIO VMware environment and performing a clean install	266
• Configuring ScaleIO devices in Linux LVM	267
• Configuring session timeout parameters	268
• Fixing keytool errors	268
• Installing Java on SUSE 12 servers	269
• SVM manual memory allocation	269
• Upgrading Java	271
• Mounting ScaleIO	271
• The ScaleIO Gateway web server isn't responding	273
• Upgrading the Gateway when a custom certificate is used	274
• Uploading a new OVA	275
• Using the same data network for different NICs	275
• What to do when the default self-signed certificate expires	275
• Add another IP address subnet to an MDM cluster	275
• Shutdown or restart a node gracefully	277
• Deployment of ScaleIO using a non-root user	280

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 HardiskX` SCSI query command.
The result of this command includes the EMC volume ID at the bottom of the output.
2. On the MDM, get the ScaleIO volume information:

```
C:\Program Files\emc\scaleio\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
```

- 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

- 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 18 Default ports

Port used by	Port #	Protocol	File to change	Field to modify (or to add, if it does not exist)	Notes
MDM listener	6611	Protobuf over TCP	Note Cannot be modified, and must be available		
MDM Cluster member	9011	Protobuf over TCP	/opt/emc/scaleio/mdm/cfg/conf.txt	actor_cluster_port=<NEW_PORT>	
SDS listener	7072	Proprietary protocol over TCP	/opt/emc/scaleio/sds/cfg/conf.txt	tgt_port=<NEW_PORT>	SDCs connect through this port for data communication and to the MDM for meta-data communication. When multiple SDSs are installed on the same physical server, use ports 7072+x, where x is the index of the SDS (for example, 70721, 70722).
LIA listener	9099	Protobuf over TCP	/opt/emc/scaleio/lia/cfg/conf.txt	lia_port=<NEW_PORT>	The Installation Manager connects to the LIA to perform installation and maintenance-related operations.
Gateway-Installation Manager/REST (not secure)	80 (or 8080, together with 8443)	REST over HTTPS	<gateway installation directory>/conf/catalina.properties	http.port=80 (or 8080)	After changing the port, you must restart the service/daemon: <ul style="list-style-type: none"> Linux: Run <code>service scaleio-gateway restart</code> Windows: Restart the EMC ScaleIO Gateway service
Gateway-Installation	443 (or 8443,	REST over HTTPS	<gateway installation	ssl.port=443 (or 8443)	

Table 18 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
```

- Ubuntu/hLinux. These files must be extracted before use, as described in the "Manual deployment on physical servers" section of the *ScaleIO Deployment Guide*.

```
MDM_IP=<LIST_VIP_MDM_IPS> dpkg -i <SDC_PATH>.deb
```

- CoreOS

```
MDM_IP=<LIST_VIP_MDM_IPS> ./<LIST_VIP_MDM_IPS>.bsx
```

where

- *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM
- *<SDC_PATH>* is the path where the SDC installation package is located

Results

The SDC is installed on the Linux server and is connected to the ScaleIO system.

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/.*/" ]
```

- 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
```

- 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`
- 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.
 - 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.
 - After changing the parameters, restart the MDM service (delete and create service) for the changes to take effect.
 - 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

- For SVMs that are SDS-only, perform the following:
 - Move the SDS to maintenance mode (MM).
 - Shut down the SVM.
 - Increase SVM memory, according to the formula below.
 - Power up the SVM.
 - Exit MM.
- For SVMs that are MDM (Master, Slave, or TB, may contain SDS, also):
 - Start with Slaves and TBs:
 - Move the SDS to maintenance mode (MM).
 - Shut down the SVM.
 - Increase SVM memory, according to the formula below.
 - Power up the SVM.
 - Exit MM.
 - Proceed with the Master MDM:
 - Switch ownership, so the Master MDM is now a Slave MDM.
 - Move the SDS to maintenance mode (MM).
 - Shut down the SVM.
 - Increase SVM memory, according to the formula below.
 - Power up the SVM.
 - Exit MM.

The memory allocation formula:

Component	Memory allocation rules
Base SVM	<ul style="list-style-type: none"> 350 MB
MDM (Master/Slave)	<ul style="list-style-type: none"> $470 \text{ MB} + (500 \text{ KB} * 8 \text{ TB of volume capacity}) + (1.44 \text{ KB} * \text{number of volumes}) + (4 \text{ KB} * \text{number of SDS devices})$ Maximum supported volumes: 256 K
Tie Breaker MDM	<ul style="list-style-type: none"> 50 MB
SDS	<ul style="list-style-type: none"> $(\text{Base}) 536 \text{ MB} + (\text{RmCache Size}) * 1.15 + (\text{Storage capacity in TB}) * 53 \text{ MB}$ For SDS high performance profile, we add 195 MB.
SDC	<ul style="list-style-type: none"> $132 \text{ KB} + 23 \text{ MB} * (\text{number of MDMs}) + 25 \text{ KB} * (\text{number of SDSs}) + 1.5 \text{ KB} * (\text{number of volumes}) + 16 \text{ B} * (\text{number of volume blocks}) + 24 \text{ KB} * (8 \text{ TB of volume capacity})$

Component	Memory allocation rules		
	<ul style="list-style-type: none"> Volume blocks: 1 GB storage = 8 volume blocks 		
RFcache	<ul style="list-style-type: none"> 16 * (cache_size/page_size) Commonly-used sizes: 		
	RFcache page size	RFcache memory requirement, if the cache device is 800 GB	RFcache memory requirement, if the cache device is 1.6 TB
	64 K	200 MB	400 MB
	32 K	400 MB	800 MB
	16K	800 MB	1.6 GB
	8 K	1.6 GB	3.2 GB
	4 K	3.2 GB	6.4 GB

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 258.

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	C:\Program Files\EMC\ScaleIO\Gateway\webapps\ROOT\WEB-INF\classes\
Linux	/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes

To reset the Scaleio-Gateway password, perform the following steps:

Procedure

1. Use `SioGWTool` to reset the password by typing the following command:

```
SioGWTool --reset_password --password <new_scaleio-gateway_password> --config_file
<path_to_file_gatewayUser.properties>
```

Note

The path to `SioGWTool` is:

Linux: `/opt/emc/scaleio/gateway/bin/SioGWTool.sh`

Windows: `C:\Program Files\EMC\ScaleIO\Gateway\bin\SioGWTool.bat`

2. Restart the `scaleio-gateway` service

The ScaleIO Gateway web server isn't responsive and the following error appears in the catalina log file:

- Windows:
C:\Program Files\EMC\ScaleIO\Gateway\logs\catalina.<date>.log
- Linux:

```
/opt/emc/scaleio/gateway/logs
```

```
2014-06-21 22:50:57,113 [main] ERROR
o.a.coyote.http11.Http11NioProtocol - Failed to initialize end
point associated with ProtocolHandler ["http-nio-443"]
java.net.BindException: Address already in use: bind
```

Solution

Perform one of the following:

Procedure

1. Find the service/daemon that is currently occupying that port and stop it:

- Windows

Run: `netstat -anb`

- Linux

Run: `netstat -alp`

On Windows, one of the common applications that occupies this port is the VMware workstation, which uses this port for the shared VM feature. You can configure VMware workstation to use a different port via the Settings dialog, or you can disable the shared VM feature.

Once the port is free, restart the scaleio-gateway service:

- Windows

Restart the EMC ScaleIO Gateway service.

- Linux

Type the command `service scaleio-gateway restart`

2. Change the ScaleIO Gateway web server to run on a different port, as described in [“Changing default ports”](#).

After doing so, restart the ScaleIO Gateway service/daemon, as described above. Access the Gateway with the new port. For example: `https://<host>:<port>`

Upgrading the Gateway when a custom certificate is used

If a custom security certificate is used on the ScaleIO Gateway (Windows and Linux environments), you must save a copy of the certificate (`*.keystore` file) and the `catalina.properties` file before you upgrade the gateway. After the upgrade is complete, you must copy these files back to their original location.

The default file locations, per operating system, are:

Linux:

```
/opt/emc/scaleio/gateway/conf/catalina.properties
```

```
/opt/emc/scaleio/gateway/conf/certificates/.keystore
```

Windows (64 bit):

```
C:\Program Files\EMC\ScaleIO\Gateway\conf\catalina.properties
```

```
C:\Program Files\EMC\ScaleIO\Gateway\conf\certificates
\.keystore
```

Uploading a new OVA

If you have already used the OVA to create a template, you cannot create another template with the same name in the same datacenter.

Either remove the original template first, or use the `ScaleIOPluginSetup-2.5-<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".

GLOSSARY

A

- Active Directory** Active Directory (AD) provides directory-based identity-related services. It maintains a directory that is used to centrally store identity information and security principles, and uses them to authenticate and authorize users and devices.
- Active Forward Rebuild** A copy of stored data is currently being rebuilt on another server, due to planned or unplanned shutdown of a server.

B

- Backward Rebuild** Data is rebuilt on servers that went offline and became active again. Forward rebuilds can take a long time, and therefore, it can be quicker to restore and update the data on a server which has come back online, than it is to do an entire rebuild on a different server.
- BWC** Bandwidth counters.

C

- Cache** Cache is random access electronic storage used to retain frequently used data for faster access by the channel. Cache is a critical aspect of storage performance. ScaleIO uses server DRAM for Read RAM Cache (RMcache) as well as SSD/Flash devices (RFcache) for caching reads. ScaleIO cache uses recently-accessed (LRU) data readily available to manage caching. I/Os read from cache have a lower response time than I/Os serviced by the drives. In addition, cached I/Os reduce the data drive workload, which in many cases is a performance bottleneck in the system.
- CacheCade** Read and Write caching of storage devices performed by one or more designated SSD devices in the ScaleIO system.
- Cache Hit Rate** The percentage of I/Os from cache.
- Cache Skip** Data is written directly to storage, bypassing the cache. Reasons for cache skips include: I/Os were too large, the cache device was busy, or I/Os were unaligned. The cache can also be configured to always work in passthrough mode.
- Cache Writes Handling Mode** The caching write-mode used by the system: passthrough mode (writes to storage only), or cached mode (by default, writes both to cache and to storage).
- Cluster Mode** ScaleIO is controlled by a cluster of MDM nodes, minimally consisting of a Master MDM, Slave MDM, and a Tie Breaker node. 5-node clusters consist of one Master MDM, two Slave MDMs, and two Tie Breakers.

D

Degraded Capacity	The capacity is available, but is not protected in case of another failure
Device	Physical storage device, such as a flash drive, or magnetic disk
DirectPath	In ScaleIO documentation, we use the term DirectPath to refer to the VMware vSphere VMDirectPath I/O feature.
DRL	Dirty Region Logging: DRL bits indicate if data is in-writing to a certain location. Once the data is written in both primary and secondary locations, the DRL bit associated with the written location is cleared. These bits can be either stored in DRAM only (memory_only) or also backed up in non-volatile memory (hardened). The former delivers better I/O performance; the latter reduces data movement following a power-cycle giving rise to a faster rebuild.

F

Failed Capacity	The capacity is inaccessible due to a failure, and data integrity is at risk
Fault Sets	A logical entity that ensures that SDS data is backed up on SDSs that belong to other Fault Sets, thus preventing double-point-of-failure scenarios if rack power outages occur.
Forward Rebuild	Data in storage will be rebuilt on another server, due to planned or unplanned shutdown of a server.

I

ID	Identifier, a unique sequence of characters that identifies an object in the system. In some CLI commands, an ID can be used to specify a system component.
IP Role	The role of the IP address configured for an SDS. Each SDS can have several IP addresses associated with it. Each IP address can serve a different purpose, or role. IP roles include: SDS, SDC, or both SDS and SDC.

L

LDAP	The Lightweight Directory Access Protocol (LDAP) is a directory service protocol that runs on a layer above the TCP/IP stack. It provides a mechanism used to connect to, search, and modify Internet directories using Client-Server architecture. In ScaleIO, LDAP is the protocol used by the MDM to communicate with Active Directory (AD) for authentication purposes.
Lockbox	Lockbox is a component of the RSA Common Security Toolkit (CST) which securely stores data (such as passwords) in an encrypted file. A lockbox must be defined for LDAP (secure LDAP), SNMP, and ESRS. For LDAP, lockbox use is optional.

M

Management IPs	The IP addresses of the MDMs defined in the system that can be used to access the MDM from CLI, GUI and REST.
Management Port	The Port number used by the MDM for purposes of communicating with the nodes in the ScaleIO network.
Manager MDM	An MDM that can act as a Master or a Slave in the cluster. Manager MDMs have a unique system ID, and can be given unique names. A manager can be a standby or a member of the cluster.
Master MDM	The MDM in the cluster that controls the SDSs and SDCs.
MDM	Any server with the MDM package installed on it. An MDM can be given a Manager or a Tie Breaker (default) role, during installation. MDMs have a unique MDM ID, and can be given unique names.

P

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

R

RAM Read Cache (RMcache)	Server RAM that is reserved for caching storage devices in a Storage Pool.
Read Flash Cache (RFcache)	Read-only caching of storage devices performed by one or more designated SSD devices and PCIe flash devices in a ScaleIO system.
Rebalance	When ScaleIO detects lopsided use of storage capacity, or when new nodes are added, it redistributes data across the nodes, in order to improve performance.
Rebuild	When ScaleIO detects a failure in the network, it creates a new copy of the data from the failed component, in a new location, to ensure data integrity.
Restricted MDM Mode	A mode set in which commands can only be performed from an MDM machine.
Restricted SDC Mode	Only approved SDCs can access the MDM. When this mode is enabled, volumes can only be added to approved SDCs.

S

- SDBG** The ScaleIO Debugger is a ScaleIO tech support troubleshooting tool, used to investigate for "live" systems that retrieves internal information from different ScaleIO components.
- SDC** ScaleIO Data Client, a lightweight device driver that exposes ScaleIO volumes as block devices to the application residing on the same server on which the SDC is installed.
- SDS** ScaleIO Data Server, which manages the capacity of a single server and acts as a back-end for data access. The SDS is installed on all servers contributing storage devices to the ScaleIO system.

Secondary MDM See [Slave MDM](#).

Single Mode A single MDM manages the ScaleIO network. This mode has no backup protection, and should not be used in production environments.

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

Snapshot Capacity The amount of capacity occupied by snapshots of volumes.

Spare Capacity Capacity that is reserved for system use, when recovery from failure is required. This capacity cannot be used for storage purposes.

Spare Percentage Policy This policy determines the amount of capacity that must always be reserved as free space.

Standby MDM An MDM node that is ready to use, with an ID, that has been locked to a specific ScaleIO system.

Storage Pool A sub-set of physical storage devices in a Protection Domain. Each storage device can only belong to one Storage Pool. User volumes will always use the storage of a single Storage Pool.

T

Thick Capacity Capacity allocated for thick volumes.

Thick Provisioned Volume In virtual storage, thick provisioning is a type of storage allocation in which the amount of storage capacity on a disk is pre-allocated on physical storage at the time the disk is created, meaning that the volume has all its capacity pre-allocated on creation.

Thin Capacity Capacity allocated for thin volumes.

Thin Provisioned Volume Thin provisioning is a method of optimizing the efficiency with which the available space is utilized in storage area networks (SAN). Thin provisioning operates by allocating disk storage space in a flexible manner among multiple users, based on the minimum space required by each user at any given time.

Throttling Throttling controls resource prioritization for rebuild and rebalance processes. Throttling can be controlled per Protection Domain or per Storage Pool (by configuring rebuild and rebalance policies).

Tie Breaker The Tie Breaker (TB) is an MDM that does not have a manager role, whose sole purpose is to help determine which MDM module is the manager that will become the master MDM and take control over the ScaleIO cluster.

The Tie Breaker ensures that there will always be one Master MDM achieving cluster quorum. In a 3-node cluster, there is one TB; in a 5-node cluster, there are two TBs.

U

Unavailable Capacity Capacity that is not being used, but is also unavailable (due to server outage).

Unused Capacity Capacity that is not currently being used for any purpose in the system.

V

Volume A general term referring to a storage device. In the ScaleIO system, a volume consists of multiple blocks spread evenly on Storage Pool devices.

W

Widget The full screen view can be minimized into a widget, which is a small window that floats on your screen, over other applications. Property sheets can also be minimized into widgets.

Write Misses Write requests that were not found in cache

