EMC® ScaleIO®

Version 2.0.x

Upgrade Guide

P/N 302-003-334

REV 04

EMC²

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

Published March 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.

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

CONTENTS

Figures 9

Tables 11

Preface 13

Part 1 Introduction to Upgrading 15

Chapter 1 Introduction 17
Automated upgrade of software components 18
Manual upgrade of software components 19

Part 2 Automated Upgrade 21

Chapter 2 Automated Upgrade of Linux Servers 23
Automated upgrade of physical Linux configurations 24
Linux package names 24
Preparing to upgrade physical Linux configurations 25
Upgrading from v1.32.x with the IM - Linux server, LIA installed 26
After upgrading 27
Upgrading from v1.32.x with the IM - Linux server, no LIA installed 27
After upgrading 30
Upgrading from v1.32.5 with the IM - XenServer 6.5 30
Upgrading from v2.0 or v2.0.0.1 with the IM - Linux server 32
Upgrading from v2.0.0.2+ with the IM - Linux server 34
After upgrading 36

Chapter 3 Automated Upgrade of Windows Servers 37
Upgrading physical Windows configurations 38
Preparing to upgrade physical Windows configurations 38
Upgrading from v1.32.x with the IM - Windows server, LIA installed 39
After upgrading 40
Upgrading from v1.32.x with the IM - Windows server, no LIA installed 40
After upgrading 42
Upgrading from v2.0 or v2.0.0.1 with the IM - Windows server 42
Upgrading from v2.0.0.2+ with the IM - Windows server 43

Chapter 4 Automated Upgrade of ESXi Servers 47
Upgrading ESXi configurations 48
Preparing to upgrade ESXi configurations 48
Upgrading from ScaleIO v1.32.x - ESXi server 49
Unregister the old ScaleIO plug-in 49
Register the new ScaleIO plug-in 50
Register the existing system 50

Create the datastore templates..... 51
Upgrade the Gateway.....52
Upgrade the ScaleIO components..... 52
Upgrade and restart the SDCs.....53
Upgrading from ScaleIO v2.0 or v2.0.0.1 - ESXi server..... 54
Unregister the old ScaleIO plug-in..... 54
Register the new ScaleIO plug-in..... 55
Register the existing system..... 55
Upgrade the Gateway.....55
Upgrade the ScaleIO components..... 56
Upgrading from ScaleIO v2.0.0.2+ - ESXi server.....57
Unregister the old ScaleIO plug-in..... 57
Register the new ScaleIO plug-in.....57
Register the existing system.....58
Create the datastore templates.....58
Upgrade the Gateway.....59
Upgrade the ScaleIO components..... 60
Upgrade and restart the SDCs..... 61

Chapter 5 Automated Upgrade of Servers in a 2-Layer Configuration 63
Preparing to upgrade 2-layer configuration - Linux..... 64
Upgrading 2-layer configurations (Linux and ESX).....64
Upgrading v1.32.x 2-layer (Linux) with LIA installed..... 64
2-layer (Linux)—v1.32.x, no LIA installed..... 66
Upgrading v2.0 or v2.0.0.1 2-layer.....69
Upgrading v2.0.0.2+ 2-layer - Linux server..... 71

Part 3 Manual Upgrade 73

Chapter 6 Manual Upgrade of Linux Server 75
Manual upgrade of physical Linux configurations.....76
Linux package names.....76
Manual upgrade of Linux system from v1.32.x..... 77
After upgrading..... 82
Manual upgrade of Linux system from v2.0 or v2.0.0.1, in a 3-node cluster.... 83
After upgrading..... 83
Manual upgrade of Linux system from v2.0 or v2.0.0.1, in a 5-node cluster.... 83
After upgrading..... 83
Manual upgrade of Linux system from v2.0.0.2+, in a 3-node cluster..... 84
After upgrading..... 89
Manual upgrade of Linux systems from v2.0.0.2+, in a 5-node cluster..... 89
After upgrading..... 94

Chapter 7 Manual Upgrade of Windows Server 95
Manual upgrade of physical Windows configurations.....96
Manual upgrade of Windows server from v1.32.x..... 96
After upgrading.....101
Manual upgrade of Windows server from v2.0 or v2.0.0.1, in a 3-node cluster..... 101
Manual upgrade of Windows server from v2.0 or v2.0.0.1, in a 5-node cluster..... 101

Manual upgrade of Windows server from v2.0.0.2+, in a 3-node cluster... 102
Manual upgrade of Windows server from v2.0.0.2+, in a 5-node cluster... 106

Chapter 8 Manual Upgrade of Xen Server 111
Manual upgrade of Xen system from v1.32.x....112
After upgrading....117

Part 4 Firmware Upgrade 119

Chapter 9 ScaleIO Ready Node Server Firmware Upgrades 121
Single server firmware and BIOS upgrade to a specific version....122
Upgrading the Dell firmware and BIOS.... 122
Preparing the node for a graceful reboot.... 125
Return the node to operation.... 126
Multiple firmware upgrades on one or more servers....127
Opening the KVM console.... 127
Updating the BIOS, firmware and settings.... 128

Chapter 10 Upgrade of LSI RAID Controller Firmware and Driver 133
LSI RAID controller upgrade on ESX.... 134
Preparation.... 134
Uploading firmware, driver and storcli installers.... 134
Upgrading the driver.... 135
Upgrading the firmware.... 136
Post-upgrade tasks.... 138
LSI RAID controller upgrade on Linux....138
Preparation.... 138
Uploading firmware, driver and storcli installers.... 139
Upgrading the driver.... 139
Upgrading the firmware.... 140
Post-upgrade tasks.... 143

Part 5 Reference Topics 145

Chapter 11 Post-Upgrade and Other Related Activities 147
System analysis overview....148
Creating the system analysis report.... 149
System analysis report description.... 151
System requirements....155
ScaleIO cluster components....155
Physical server requirements.... 156
Supported operating systems....157
GUI server requirements.... 158
ScaleIO Gateway server requirements.... 159
Other requirements....159
Extracting ScaleIO packages....160
Upgrading the ScaleIO GUI.... 160
Upgrading the SDC.... 160
Manually upgrading the SDC version in ESX environment.... 161
Modifying SDC parameters 162
Non-disruptive SNMP upgrade issues....163

Installing with the full Installation Manager..... 164
Installing with the Installation Manager..... 164
Configuring Installation Manager properties.....170
Upgrading the Gateway when a custom certificate is used..... 171
Upgrading the BMC firmware on a VxRack Node 100 Series node..... 171
Unregistering the ScaleIO plug-in.....172
Registering the ScaleIO plug-in manually..... 172
Installing the SDC on ESX hosts.....174
Upgrading Windows servers when the OS is installed on SATADOM.....176
Troubleshooting plug-in registration issues..... 176
SVM manual memory allocation.....177
Manual migration of an RDM-based ScaleIO system to DirectPath-based.... 178
Deploying on Ubuntu, Oracle Linux, or CoreOS servers..... 180
Ensuring the kernel version is correct..... 180
Creating the configuration file via the ScaleIO Gateway..... 181
Creating the configuration file manually..... 183
Creating a mirror repository..... 185
Maintaining a ScaleIO system..... 186
Installing RFcache on servers in an existing ScaleIO system..... 186
Extending the MDM cluster from 3 to 5-node..... 187
Creating a Lockbox for SNMP, ESRS, or LDAP.....190
Switching to secured authentication mode..... 190
Using SCLI in non-secure mode.....194
Extending an existing ScaleIO system..... 194
Configuring virtual IP addresses using Installation Manager..... 195
Configuring virtual IP addresses..... 196
Removing ScaleIO..... 197
Configure ESRS after upgrading..... 198
Enabling ESRS certificate verification.....199
Adding the ESRS Gateway's certificate to the truststore..... 199

Chapter 12 Troubleshooting 201
Troubleshooting ScaleIO..... 202
After SDC installation, the ScaleIO SVM does not start automatically.....202
Application server does not see a ScaleIO volume..... 202
Cannot log in to the Installation Manager after upgrade..... 203
Certificate error when installing SDC on Windows servers..... 204
Enabling acceleration in a new node..... 204
Gateway server recovery.....207
Recovering Gateway on a new server during upgrade.....207
Recovering IM with LIA trusted IP feature enabled.....209
SNMP configuration recovery post gateway crash during upgrade... 210
Installation Manager returns an error.....211
Installation with the Installation Manager fails.....212
LIA upgrade fails in Windows 2008.....212
Older SDCs cannot communicate with newer SDSs after upgrade to v2.x 212
Removing RFcache from Windows servers..... 213
Removing RFcache leftovers from the Windows OS registry..... 213
Replacing a faulty caching device in ScaleIO.....213
Replacing a faulty accelerated storage device in ScaleIO.....215
SVM manual memory allocation..... 216
RFcache is not installed on an upgraded Windows server..... 218

SDS with RFcache remains in "Disconnected, Decoupled" status after reboot..... 219
After retrying upgrade, "Failed to query node" message is displayed..... 219
After upgrading, a "Could not switch to cluster mode" error message is displayed..... 220
MDM package is not updated during upgrade.....220
Upgrade from v1.32.x causes Data Unavailability status..... 220
An SDS performance profile may not get set to High during upgrade..... 221

Chapter 13 Frequently Asked Questions 223
Install the ScaleIO GUI.....224
Associating ScaleIO volumes with physical disks..... 224
Volume information - Linux.....224
Volume information - Windows.....225
Changing default ports..... 226
Adding an external SDC to an existing ScaleIO system.....227
Installing SDC on an ESX server and connecting it to ScaleIO..... 227
Installing SDC on a Linux server and connecting it to ScaleIO..... 228
Installing SDC on a Windows server and connecting it to ScaleIO.... 229
Changing the LIA configuration file.....229
Cleaning the ScaleIO VMware environment and performing a clean install.... 230
Configuring ScaleIO devices in Linux LVM.....232
Configuring session timeout parameters.....232
Fixing keytool errors..... 233
Error during rpm installation command..... 233
Error during rpm upgrade command..... 233
Installing Java on SUSE 12 servers..... 233
Mounting ScaleIO.....234
The ScaleIO Gateway web server isn't responding.....235
The ScaleIO Gateway (REST service, Installation Manager) may be disabled:..... 235
The ScaleIO Gateway web server isn't responsive and the following error appears in the catalina log file:..... 236
Upgrading the Gateway when a custom certificate is used.....237
Uploading a new OVA..... 237
Using the same data network for different NICs.....237
What to do when the default self-signed certificate expires..... 237
Add another IP address subnet to an MDM cluster.....238

Chapter 14 DTK - Hardware Update Bootable ISO 241
Dell OpenManage DRAC Tools (RACADM).....242
Updating the hardware using remote RACADM.....242
SATADOM firmware update..... 244
Recommended BIOS and firmware settings..... 245
ScaleIO ID module..... 245
BIOS and firmware..... 245
Configuration settings..... 246
Troubleshooting the Hardware ISO.....248
General iDRAC failures..... 248
iDRAC virtual console issues.....248
iDRAC virtual media issues..... 249
Bootable ISO logs..... 249

Additional resources.....250
ScaleIO resources.....250
Dell Lifecycle Controller (LC).....251
Dell OpenManage Deployment Toolkit (DTK).....251
Setting up the BMC (iDRAC) IP address and BIOS.....251
Verifying status of system hardware, storage controller, and disks.....253

Glossary **255**

FIGURES

1 Table view toggle options....125
2 SDS displayed in maintenance mode....126
3 Table view toggle options....129
4 SDS displayed in maintenance mode....130
5 Before extending....188
6 After extending....188
7 Set Virtual IPs for ScaleIO system screen....196
8 Configure virtual IPs dialog box....197

TABLES

1 Upgrade support matrix..... 18
2 Upgrade support matrix..... 19
3 Linux package formats.....24
4 Linux package formats.....76
5 Server physical requirements.....156
6 Supported operating systems - ScaleIO components.....157
7 driver_sync.conf parameters..... 183
8 Default ports..... 226
9 Hardware ISO contents..... 246
10 Hardware ISO configuration settings.....246

Preface

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

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

Note

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

Related documentation

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

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

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

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

Typographical conventions

EMC uses the following type style conventions in this document:

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
Italic	Used for full titles of publications referenced in text
`Monospace`	Used for: • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
Monospace italic	Used for variables
`Monospace bold`	Used for user input
[ ]	Square brackets enclose optional values

|	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

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

Product information

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

Technical support

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

Your comments

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

PART 1

Introduction to Upgrading

This document describes how to upgrade ScaleIO software and hardware components. This part includes the following chapters:

Chapter 1, "Introduction"

CHAPTER 1

Introduction

This chapter introduces the upgrade procedure for ScaleIO components. Topics include:

- Automated upgrade of software components 18
- Manual upgrade of software components 19

Automated upgrade of software components

You can upgrade ScaleIO software components with automated tools, in a non-disruptive manner. The following table lists the versions from which the upgrade can be performed:

Table 1 Upgrade support matrix

Component	Version supported
Backend (MDM and SDS)	1.32.4 or later
Frontend (SDC)	1.32.0 or later

NOTICE

If RFcache is in use in the ScaleIO system to be upgraded, you must disable the RFcache service and remove RFcache devices from ScaleIO before upgrading.

After upgrading, add the RFcache devices and then enable the RFcache service.

To upgrade from earlier versions, contact EMC Support.

Before upgrading, ensure that the following issues are addressed:

- All servers meet the system requirements, described in System requirements on page 155.
- Ensure that there are no failed disks (failed capacity).
- Verify that the cluster state is not degraded.
- Ensure that all data NICS are connected.

It is highly recommended to perform upgrades using the Installation Manager or the VMware plug-in. Manual upgrades are described in Manual upgrade of software components on page 19.

Note

To upgrade the replication splitter for RecoverPoint, refer to the manual upgrade procedure in the *ScaleIO Write Splitter for RecoverPoint Technical Notes*. If you are using the IM to upgrade nodes that work with the splitter for RecoverPoint, you should exclude the RPA nodes from the upgrade, as described in the *ScaleIO Deployment Guide*. Replication support is version-specific. Refer to the ESSM for complete information.

When using the IM to upgrade, you can configure the following settings as part of the upgrade process:

- Set high performance profile

  Choose to set the high-performance profile, instead of the default profile, for MDM, SDS, or SDC components. Before upgrading multiple SDS servers and changing their performance profile from default to high-performance, ensure that there is adequate available RAM on each server. If there isn't, you can change the performance profile after upgrading, one server at-a-time.
- Restart Windows servers

  Enable the automatic restart of Windows servers on which SDC or RFcache are being installed.

- Enable parallel SDS upgrades

  Enable the upgrading of all SDS servers in an entire Fault Set at one time.

These upgrade procedures leave the system in non-secure mode. After upgrading, you can configure the following settings:

- Secure communication mode

  Note

  If you choose not to switch to secure mode, you need to disable secure communication on every MDM server before you can run SCLI commands.

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses (not for Windows servers)

These options, and others, are described in "Maintaining a ScaleIO system".

The upgrade procedure for your environment is determined by your current configuration. Proceed to the section that matches your configuration:

Configuration	See page
Physical Linux - The base system resides solely on bare metal Linux servers.	Automated upgrade of physical Linux configurations on page 24
Physical Windows - The base system resides solely on bare metal Windows servers.	Upgrading physical Windows configurations on page 38
ESXi - The base system resides solely on bare metal ESXi servers.	Upgrading ESXi configurations on page 48
2-layer (ESX and Linux) - The MDM and SDS reside on bare metal Linux servers, and the SDCs reside on ESXi servers.	Upgrading 2-layer configurations (Linux and ESX) on page 64

Manual upgrade of software components

You can upgrade ScaleIO software components manually, in a non-disruptive manner. The following table lists the versions from which the upgrade can be performed:

Table 2 Upgrade support matrix

Component	Version supported
Backend (MDM and SDS)	1.32.0 or later
Frontend (SDC)	1.32.0 or later

NOTICE

If RFcache is in use in the ScaleIO system to be upgraded, you must disable the RFcache service and remove RFcache devices from ScaleIO before upgrading.

After upgrading, add the RFcache devices and then enable the RFcache service.

To upgrade from earlier versions, contact EMC Support.

It is highly recommended to perform upgrades using the Installation Manager or the VMware plug-in.

Note

When performing a manual upgrade (as opposed to using the Installation Manager), it is crucial to follow the described upgrade procedures step by step. Failing to do so might result in data loss and system instability.

Before upgrading, ensure that the following issues are addressed:

- All servers meet the system requirements, described in System requirements on page 155.
- Ensure that there are no failed disks (failed capacity).
- Verify that the cluster state is not degraded.
- Ensure that all data NICS are connected.

Run the ScaleIO system analysis to analyze the ScaleIO system immediately prior to upgrading. It will identify and avoid faulty setups in the system. For more information, see System analysis overview on page 148.

Note

To upgrade the replication splitter for RecoverPoint, refer to the manual upgrade procedure in the *ScaleIO Write Splitter for RecoverPoint Technical Notes*. Replication support is version-specific. Refer to the ESSM for complete information.

After upgrading, you can configure the following settings:

- Secure communication mode

  Note

  If you choose not to switch to secure mode, you need to disable secure communication on every MDM server before you can run SCLI commands.

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses (not for Windows servers)

These options, and others, are described in "Maintaining a ScaleIO system".

The upgrade procedure for your environment is determined by your current configuration. Proceed to the section that matches your configuration:

Configuration	See page
Physical Linux - The base system resides solely on bare metal Linux servers.	Manual upgrade of physical Linux configurations on page 76
Physical Windows - The base system resides solely on bare metal Windows servers.	Manual upgrade of physical Windows configurations on page 96
ESXi and 2-layer (ESX and Linux)	Contact EMC Customer Support
Physical Xen - The base system resides solely on bare metal Xen servers.	Manual upgrade of Xen system from v1.32.x on page 112

PART 2

Automated Upgrade

This part describes how to upgrade ScaleIO components using ScaleIO automated management tools. Chapters include:

Chapter 2, "Automated Upgrade of Linux Servers"

Chapter 3, "Automated Upgrade of Windows Servers"

Chapter 4, "Automated Upgrade of ESXi Servers"

Chapter 5, "Automated Upgrade of Servers in a 2-Layer Configuration"

CHAPTER 2

Automated Upgrade of Linux Servers

This chapter describes how to use the ScaleIO automated management tools to upgrade ScaleIO components on Linux servers. Topics include:

- Automated upgrade of physical Linux configurations 24
- Preparing to upgrade physical Linux configurations 25
- Upgrading from v1.32.x with the IM - Linux server, LIA installed 26
- Upgrading from v1.32.x with the IM - Linux server, no LIA installed 27
- Upgrading from v1.32.5 with the IM - XenServer 6.5 30
- Upgrading from v2.0 or v2.0.0.1 with the IM - Linux server 32
- Upgrading from v2.0.0.2+ with the IM - Linux server 34

Automated upgrade of physical Linux configurations

This section describes how to upgrade when all ScaleIO components reside on physical Linux servers.

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.

Note

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 already installed, you will first upgrade it to the new version. If no LIA is installed, you will install the new version.

Upgrading is non-disruptive; you can upgrade while IOs are running and volumes are mapped.

You can use the IM to upgrade from a 3-node cluster to a 3-node cluster, or from a 5-node to a 5-node. Extending from a 3-node to a 5-node cluster is performed (and described) after upgrading.

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>-2.0-13000.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 3 Linux package formats

Linux flavor	Package format	Example
CoreOS	CoreOS	`EMC-ScaleIO-mdm-2.0-13000.X.CoreOS.x86_64.tar` CoreOS packages may need to be extracted before use. This is described where relevant.
RHEL/OL/CentOS	el<version>	`EMC-ScaleIO-mdm-2.0-13000.X.el<version>.x86_64.rpm` Example: `EMC-ScaleIO-mdm-2.0-13000.X.el6.x86_64.rpm`
SUSE	sles<version>	`EMC-ScaleIO-mdm-2.0-13000.X.sles<version>.x86_64.rpm` Example: `EMC-ScaleIO-mdm-2.0-13000.X.sles11.3.x86_64.rpm`
Ubuntu	Ubuntu.<version>	`EMC-ScaleIO-mdm-2.0-13000.X.Ubuntu.<version>.x86_64.tar` Example: `EMC-ScaleIO-mdm-2.0-13000.X.Ubuntu.16.04.x86_64.tar`

Table 3 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-2.0-13000.X.xs<version>.x86_64.rpm` Example: `EMC-ScaleIO-mdm-2.0-13000.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 upgrading CoreOs, Oracle Linux (OL), or Ubuntu systems, ensure that the ScaleIO environment is prepared, as described in Deploying on Ubuntu, Oracle Linux, or CoreOS servers on page 180.

Preparing to upgrade physical Linux configurations

Before you begin

Run the ScaleIO system analysis to analyze the ScaleIO system immediately prior to upgrading. It will identify and avoid faulty setups in the system. For more information, see System analysis overview on page 148.

Before upgrading the ScaleIO components, perform the following steps:

Procedure

1. Install Java 1.8 on the Gateway and GUI servers. For example:
   - Linux:

     ```
     rpm -i /var/tmp/jre-8uXX-linux-x64.rpm
     ```

   - Windows:

     ```
     jre-8uXX-windows-x64.exe
     ```

2. Download and extract the ScaleIO packages to `/tmp/`
3. Upgrade the GUI:

   ```
   EMC-ScaleIO-gui-2.0-13000.X.msi
   ```

After you finish

The next upgrade steps for your environment are determined by the base version of ScaleIO and whether the LIA is installed. Proceed to the section that matches your environment.

Upgrading from v1.32.x with the IM - Linux server, LIA installed

Use the Installation Manager (IM) to upgrade ScaleIO v1.32.4 (or later) components, on a Linux server on which the LIA is installed.

To upgrade from pre-1.32.4 versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.
The following upgrade procedure can be used for upgrading a 3 or 5-node cluster on all Linux flavor servers, with the exception of CoreOS and Xen servers. To upgrade these servers running ScaleIO v1.32.x:

- CoreOS manual upgrade - see Manual upgrade of Linux system from v1.32.x on page 77.
- XenServer 6.5 and ScaleIO v1.32.5 automatic upgrade - see Upgrading from v1.32.5 with the IM - XenServer 6.5 on page 30.
- XenServer and all other ScaleIO v1.32.x versions - see Manual upgrade of Xen system from v1.32.x on page 112

When SNMP is enabled in the 1.32.x base version, the previous SNMP MDM credentials are erased from the old `gatewayUser.properties` configuration file and moved into the v2.x Lockbox. To harden the password, you can use the SioGWTool after the upgrade, or you can pass an environment variable while upgrading the Gateway. For more information, see "Non-Disruptive Upgrade Issues" in the *ScaleIO Upgrade Guide*.

Procedure

1. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

   ```
   SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
   gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
   ```

   Example:

   ```
   SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
   ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
   ```

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

2. Log in to the IM server (https:*//<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in “Installing with the Installation Manager”.

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

3. Upgrade the ScaleIO components:

a. From the IM, click the **Maintain** tab.

b. From the **Maintenance operation** screen, perform the upgrade:

   a. Type the Master (or Primary) MDM credentials.

   b. Type the LIA password.

   c. From the **Advanced Options** section, select both of the **Allow non-secure communication** options.

      Security warnings are displayed because the new Gateway is trying to communicate in a non-secure mode. You can ignore the messages.

   d. Click **Retrieve system topology**. A success message is displayed.

   e. Click **Upgrade**.

c. Set upgrade options:

   - Set performance profiles.
   - Automatic restart of Windows servers.
   - Enable parallel SDS upgrades.

d. Confirm the upgrade operation and click **Upgrade**.

   Note

   If a certificate notice is displayed, review and approve it. You may need to rerun the phase to continue.

e. Click the **Monitor** tab, and approve each step of the upgrade process.

f. When complete, click **Mark operation completed**.

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 148.

After upgrading

After the upgrade, you can perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses

These options, and others, are described in "Maintaining a ScaleIO system".

Upgrading from v1.32.x with the IM - Linux server, no LIA installed

This section describes how to upgrade from ScaleIO, from v1.32.4 (or later), when LIA is not installed. To upgrade from pre-1.32.4 versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.

The following upgrade procedure can be used for upgrading a 3 or 5-node cluster on all Linux flavor servers, with the exception of CoreOS and Xen servers. To upgrade these servers running ScaleIO v1.32.x:

- CoreOS manual upgrade - see Manual upgrade of Linux system from v1.32.x on page 77.
- XenServer 6.5 and ScaleIO v1.32.5 automatic upgrade - see Upgrading from v1.32.5 with the IM - XenServer 6.5 on page 30.
- XenServer and all other ScaleIO v1.32.x versions - see Manual upgrade of Xen system from v1.32.x on page 112

When SNMP is enabled in the 1.32.x base version, the previous SNMP MDM credentials are erased from the old `gatewayUser.properties` configuration file and moved into the v2.x Lockbox. To harden the password, you can use the SioGWTool after the upgrade, or you can pass an environment variable while upgrading the Gateway. For more information, see "Non-Disruptive Upgrade Issues" in the *ScaleIO Upgrade Guide*.
The first step is to install the LIA; then proceed to upgrade the rest of the system.

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-2.0-13000.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. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

If the Gateway isn't installed run this command:

```
GATEWAY_ADMIN_PASSWORD=<GW password> rpm -U /tmp/EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

Note

To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

---

4. Log in to the IM server (https:*//<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in "Installing with the Installation Manager".

Note

To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

---

5. Upgrade the ScaleIO components:
   a. From the IM, click the **Maintain** tab.
   b. From the **Maintenance operation** screen, perform the upgrade:
      a. Type the Master (or Primary) MDM credentials.
      b. Type the LIA password.
      c. From the **Advanced Options** section, select both of the **Allow non-secure communication** options.

         Security warnings are displayed because the new Gateway is trying to communicate in a non-secure mode. You can ignore the messages.
      d. Click **Retrieve system topology**. A success message is displayed.
      e. Click **Upgrade**.
   c. Set upgrade options:
      - Set performance profiles.
      - Automatic restart of Windows servers.
      - Enable parallel SDS upgrades.

d. Confirm the upgrade operation and click **Upgrade**.

Note

If a certificate notice is displayed, review and approve it. You may need to rerun the phase to continue.

---

e. Click the **Monitor** tab, and approve each step of the upgrade process.

f. When complete, click **Mark operation completed**.

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 148.

After upgrading

After the upgrade, you can perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses

These options, and others, are described in "Maintaining a ScaleIO system".

Upgrading from v1.32.5 with the IM - XenServer 6.5

Use the Installation Manager (IM) to upgrade ScaleIO v1.32.5 components, on a XenServer 6.5. Use this procedure for all systems, whether the LIA is installed or not.

All upgraded SDC nodes must be restarted as part of the upgrade. You will be instructed to migrate VMs after the upgrades are complete.

To upgrade from other 1.32.x versions of ScaleIO, either upgrade to v1.32.5 version, then proceed with the upgrade described here, or use the manual upgrade described in Manual upgrade of Xen system from v1.32.x on page 112.

When SNMP is enabled in the 1.32.x base version, the previous SNMP MDM credentials are erased from the old `gatewayUser.properties` configuration file and moved into the v2.x Lockbox. To harden the password, you can use the SioGWTool after the upgrade, or you can pass an environment variable while upgrading the Gateway. For more information, see "Non-Disruptive Upgrade Issues" in the *ScaleIO Upgrade Guide*.

Procedure

1. Log in to the v1.32.5 IM server (https://*<IM_Server_IP>*). If necessary, accept the certificate warning.
2. Enter the default user name, admin, and the password defined when the IM was prepared, then click **Login**.
3. From the **Welcome** screen, click the **Maintain** tab.
4. From the **Maintenance Operation** screen, upgrade the LIA:

   a. Type the Primary MDM IP address and password.

b. Select **Use native SSH/WMI**

c. Browse to and select the current deployment CSV file.

d. Click **Retrieve system topology**.

The system topology is displayed.

e. Click **Upgrade LIA**. You may need to confirm this step.

f. Click the **Monitor** tab, and approve each step of the upgrade process.

g. When the process is complete, click **Mark operation completed**.

The LIA is now upgraded on all nodes.

5. Upgrade the currently installed (or install a new) Gateway:

OS	Procedure

Linux	Run the following command (all on one line):
	`SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-gateway-2.0-13000.X.x86_64.rpm`
	Example:
	`SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-ScaleIO-gateway-2.0-13000.X.x86_64.rpm`

Windows	Run the following command: `EMC-ScaleIO-gateway-2.0-13000.X-x64.msi`

A success message is displayed, and the IM interface now reflects the updated Gateway version.

6. Log in to the IM server (https://*<IM_Server_IP>*), and upload the installation packages, as described in "Installing with the Installation Manager".

Note

To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work. The current TB package may be displayed as the latest version even after you upload the new packages. You can ignore this.

7. Upgrade the ScaleIO components:

a. From the IM, click the **Maintain** tab.

b. From the **Maintenance operation** screen, perform the upgrade:

a. Type the Master (or Primary) MDM credentials.

b. Type the LIA password.

c. From the **Advanced Options** section, select both of the **Allow non-secure communication** options.

Security warnings are displayed because the new Gateway is trying to communicate in a non-secure mode. You can ignore the messages.

d. Click **Retrieve system topology**. A success message is displayed.

e. Click **Upgrade**.

c. Set upgrade options:

- Set performance profiles.
- Enable parallel SDS upgrades.

d. Confirm the upgrade operation and click **Upgrade**.

Note

If a certificate notice is displayed, review and approve it.

---

e. Click the **Monitor** tab, and approve each step of the upgrade process.

8. Restart an SDC node:

   a. On the SDC node to be restarted, migrate the VMs (if any).

   b. Restart the SDC.

   c. Before restarting the next SDC:

      a. Ensure that the SDC has re-established communication with the MDM cluster.

      b. If the SDC was also an MDM, ensure that the cluster state has returned to normal.

      c. If the SDC was also an SDS, ensure that any rebuild/rebalance processes have completed.

      d. If the SDC is both an MDM and an SDS, ensure that all of the listed activities have completed.

9. After ensuring that all activities are complete, restart the other SDC nodes, one-at-a-time, as described in the previous step. Continue until all SDCs have been restarted.

10. When complete, click **Mark operation completed**.

Upgrading from v2.0 or v2.0.0.1 with the IM - Linux server

Use the Installation Manager (IM) to upgrade ScaleIO components on a Linux server, from v2.0 or v2.0.0.1 to the current version.

Before you begin

Before upgrading Ubuntu or CoreOS servers, you must extract the LIA package, as described in Extracting ScaleIO packages on page 160.

This procedure can be used for upgrading a 3 or 5-node cluster on all Linux flavor servers.

When SNMP is enabled in the 2.x base version, SNMP MDM credentials remain in the v2.x Lockbox after the upgrade.

This procedure is completed with the following steps:

1. Upgrade to v2.0.0.x (latest) - described here.
2. Upgrade from v2.0.0.x (latest) to the current version - described in Upgrading from v2.0.0.2+ with the IM - Linux server on page 34

Begin the upgrade:

Procedure

1. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

   Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

2. When upgrading Ubuntu v14.04 or CoreOS servers, perform the following to upgrade the LIA manually. For other Linux flavors, skip to the next step.
   a. Extract the LIA package, as described in Extracting ScaleIO packages on page 160.
   b. Upgrade the LIA:
      - Ubuntu:

```
dpkg -i <path to the new Ubuntu LIA package>
```

      - CoreOS:

```
./<path to the new CoreOS LIA package>
```

3. When upgrading XenServer servers, perform the following to upgrade the LIA manually. For other Linux flavors, skip to the next step.
   a. Upgrade (or install) LIA on one of the servers in the system:

```
rpm -U EMC-ScaleIO-lia-2.0-13000.X.xsX.X.X.x86_64.rpm --
nosignature
```

   b. Repeat for every server in the system.

   After all LIAs are installed/upgraded, continue to the next step.

4. Log in to the IM server (https://*<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in "Installing with the Installation Manager".

Note

To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

---

5. Upgrade the ScaleIO components:
   a. From the IM, click the **Maintain** tab.
   b. From the **Maintenance operation** screen, perform the following:
      a. Type the Master (or Primary) MDM credentials.
      b. Type the LIA password.
      c. Click **Retrieve system topology**. A success message is displayed.
      d. Click **Upgrade**.
   c. Set upgrade options:
      - Set performance profiles.
      - Automatic restart of Windows servers.
      - Enable parallel SDS upgrades.
   d. Confirm the upgrade operation and click **Upgrade**.

      Note

      If a certificate notice is displayed, review and approve it.

   e. Click the **Monitor** tab, and approve each step of the upgrade process.
   f. When complete, click **Mark operation completed**.

Results

The upgrade to v2.0.0.x (latest) is complete.

After you finish

Complete the upgrade by following the steps described in Upgrading from v2.0.0.2+ with the IM - Linux server on page 34.
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 148.

Upgrading from v2.0.0.2+ with the IM - Linux server

Use the Installation Manager (IM) to upgrade ScaleIO components on a Linux server, from v2.0.0.2 (or later) to the current version.

Before you begin

Before upgrading Ubuntu or CoreOS servers, you must extract the LIA package, as described in Extracting ScaleIO packages on page 160.

This procedure can be used for upgrading a 3 or 5-node cluster on all Linux flavor servers.
This procedure is also used as the second step for some upgrades from previous versions.

When SNMP is enabled in the 2.x base version, SNMP MDM credentials remain in the v2.x Lockbox after the upgrade.

Procedure

1. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

   Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

2. When upgrading Ubuntu v14.04 or CoreOS servers, perform the following to upgrade the LIA manually. For other Linux flavors, skip to the next step.
   a. Extract the LIA package, as described in Extracting ScaleIO packages on page 160.
   b. Upgrade the LIA:
      - Ubuntu:

```
dpkg -i <path to the new Ubuntu LIA package>
```

      - CoreOS:

```
./<path to the new CoreOS LIA package>
```

3. When upgrading XenServer servers, perform the following to upgrade the LIA manually. For other Linux flavors, skip to the next step.
   a. Upgrade (or install) LIA on one of the servers in the system:

```
rpm -U EMC-ScaleIO-lia-2.0-13000.X.xsX.X.X.x86_64.rpm --
nosignature
```

   b. Repeat for every server in the system.

      After all LIAs are installed/upgraded, continue to the next step.

4. Log in to the IM server (https://*<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in "Installing with the Installation Manager".

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

5. Upgrade the ScaleIO components:

a. From the IM, click the **Maintain** tab.

b. From the **Maintenance operation** screen, perform the following:

   a. Type the Master (or Primary) MDM credentials.

   b. Type the LIA password.

   c. Click **Retrieve system topology**. A success message is displayed.

   d. Click **Upgrade**.

c. Set upgrade options:

   - Set performance profiles.
   - Automatic restart of Windows servers.
   - Enable parallel SDS upgrades.

d. Confirm the upgrade operation and click **Upgrade**.

Note

If a certificate notice is displayed, review and approve it.

---

e. Click the **Monitor** tab, and approve each step of the upgrade process.

f. When complete, click **Mark operation completed**.

Results

The upgrade is complete.

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 148.

After upgrading

After the upgrade, you can perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses

These options, and others, are described in "Maintaining a ScaleIO system".

CHAPTER 3

Automated Upgrade of Windows Servers

This chapter describes how to use the ScaleIO automated management tools to upgrade ScaleIO components on Windows servers. Topics include:

- Upgrading physical Windows configurations 38
- Preparing to upgrade physical Windows configurations 38
- Upgrading from v1.32.x with the IM - Windows server, LIA installed 39
- Upgrading from v1.32.x with the IM - Windows server, no LIA installed 40
- Upgrading from v2.0 or v2.0.0.1 with the IM - Windows server 42
- Upgrading from v2.0.0.2+ with the IM - Windows server 43

Upgrading physical Windows configurations

This section describes how to upgrade when all ScaleIO components reside on physical Windows servers.

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.

Note

To determine if the LIA is installed, search for `EMC-scaleio-lia` in the **Control Panel** > **Programs and Features**.

If LIA is already installed, it will be upgraded as part of the described process. If no LIA is installed, you will install the new version.

Upgrading is non-disruptive; you can upgrade while IOs are running and volumes are mapped.

Extending from a 3-node to a 5-node cluster is performed, and described, after upgrading.

Preparing to upgrade physical Windows configurations

Steps to perform before upgrading ScaleIO on physical Windows servers.

When upgrading from v1.32.x, do not delete volumes or snapshots during the MDM cluster upgrade phase of the upgrade. This phase is normally very short.

Procedure

1. Install Java 1.8 on the Gateway and GUI servers:

```
jre-8u45-windows-x64.exe
```

2. Download and extract the installation files needed for your operating system.
3. Upgrade the GUI:

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

After you finish

The next upgrade steps for your environment are determined by the base version of ScaleIO and whether the LIA is installed. Proceed to the section that matches your environment:

- Upgrading from v1.32.x with the IM - Windows server, LIA installed on page 39
- Upgrading from v1.32.x with the IM - Windows server, no LIA installed on page 40
- Upgrading from v2.0 or v2.0.0.1 with the IM - Windows server on page 42
- Upgrading from v2.0.0.2+ with the IM - Windows server on page 43

Upgrading from v1.32.x with the IM - Windows server, LIA installed

This section describes how to upgrade from ScaleIO v1.32.4 (or later), when LIA is installed. To upgrade from pre-1.32.4 versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.

When SNMP is enabled in the 1.32.x base version, the previous SNMP MDM credentials are erased from the old `gatewayUser.properties` configuration file and moved into the v2.x Lockbox. To harden the password, you can use the SioGWTool after the upgrade, or you can pass an environment variable while upgrading the Gateway. For more information, see "Non-Disruptive Upgrade Issues" in the *ScaleIO Upgrade Guide*.
Before you begin, ensure that:

- The EMC ScaleIO gateway service is started.
- The Visual C++ redistributable 2010 package (64-bit) is installed on the Gateway server.

Procedure

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

   `EMC-ScaleIO-gateway-2.0-13000.X-x64.msi`

2. Run the file.

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

3. Log in to the IM server (https://*<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in "Installing with the Installation Manager".

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

4. Upgrade the ScaleIO components:

   a. From the IM, click the **Maintain** tab.

   b. From the **Maintenance operation** screen, perform the upgrade:

      a. Type the Master (or Primary) MDM credentials.

      b. Type the LIA password.

      c. From the **Advanced Options** section, select both of the **Allow non-secure communication** options.

      Security warnings are displayed because the new Gateway is trying to communicate in a non-secure mode. You can ignore the messages.

d. Click **Retrieve system topology**. A success message is displayed.

e. Click **Upgrade**.

c. Set upgrade options:

- Set performance profiles.
- Automatic restart of Windows servers.
- Enable parallel SDS upgrades.

d. Confirm the upgrade operation and click **Upgrade**.

Note

If a certificate notice is displayed, review and approve it. You may need to rerun the phase to continue.

e. Click the **Monitor** tab, and approve each step of the upgrade process.

f. When complete, click **Mark operation completed**.

After upgrading

These upgrade procedures leave the system in non-secure authentication mode. After the upgrade, you can switch to secure mode, as well as perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node

These options, and others, are described in "Maintaining a ScaleIO system".

Note

If you choose not to switch to secure mode, you need to disable secure communication on every MDM server before you can run SCLI commands.

Upgrading from v1.32.x with the IM - Windows server, no LIA installed

This section describes how to upgrade from ScaleIO, from v1.32.4 (or later), when LIA is not installed. To upgrade from pre-1.32.4 versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.

The first step is to install the LIA; then proceed to upgrade the rest of the system.

When SNMP is enabled in the 1.32.x base version, the previous SNMP MDM credentials are erased from the old `gatewayUser.properties` configuration file and moved into the v2.x Lockbox. To harden the password, you can use the SioGWTool after the upgrade, or you can pass an environment variable while upgrading the Gateway. For more information, see "Non-Disruptive Upgrade Issues" in the *ScaleIO Upgrade Guide*.
Before you begin, ensure that:

- The EMC ScaleIO gateway service is started.
- The Visual C++ redistributable 2010 package (64-bit) is installed on the Gateway server.

Procedure

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

   ```
   msiexec /i EMC-ScaleIO-lia-2.0-13000.X.msi
   TOKEN=<password>
   ```

   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 (! @#$ ...)

   The LIA component requires security configuration, as described in the *EMC ScaleIO Security Configuration Guide*.

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
   ```

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

   ```
   EMC-ScaleIO-gateway-2.0-13000.X-x64.msi
   ```

4. Run the file.

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

5. Log in to the IM server (https://<IM_Server_IP>), and for each OS in your ScaleIO environment, upload the installation packages, as described in "Installing with the Installation Manager".

After upgrading

These upgrade procedures leave the system in non-secure authentication mode. After the upgrade, you can switch to secure mode, as well as perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node

These options, and others, are described in "Maintaining a ScaleIO system".

Note

If you choose not to switch to secure mode, you need to disable secure communication on every MDM server before you can run SCLI commands.

Upgrading from v2.0 or v2.0.0.1 with the IM - Windows server

Use the Installation Manager (IM) to upgrade ScaleIO components on a Windows server, from v2.0 or v2.0.0.1 to the current version.

Before you begin

Ensure that:

- The EMC ScaleIO gateway service is started.
- The Visual C++ redistributable 2010 package (64-bit) is installed on the Gateway server.
- If a server has XCACHE (RFcache) installed on it, you must prepare the server BEFORE upgrading. The procedure is described in RFcache is not installed on an upgraded Windows server on page 218.

This procedure can be used for upgrading a 3 or 5-node cluster.

When SNMP is enabled in the 2.x base version, SNMP MDM credentials remain in the v2.x Lockbox after the upgrade.

This procedure is completed with the following steps:

1. Upgrade to v2.0.0.2 - described here.
2. Upgrade from v2.0.0.2 to the current version - described in Upgrading from v2.0.0.2+ with the IM - Windows server on page 43
   Begin the upgrade:

Procedure

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

   ```
   EMC-ScaleIO-gateway-2.0-XXXX.X-x64.msi
   ```

2. Run the file.

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

Note

To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

3. Log in to the IM server (https://*<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in "Installing with the Installation Manager".

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

4. Upgrade the ScaleIO components:
   a. From the IM, click the **Maintain** tab.
   b. From the **Maintenance operation** screen, perform the following:
      a. Type the Master (or Primary) MDM credentials.
      b. Type the LIA password.
      c. Click **Retrieve system topology**. A success message is displayed.
      d. Click **Upgrade**.
   c. Set upgrade options:
      - Set performance profiles.
      - Automatic restart of Windows servers.
      - Enable parallel SDS upgrades.
   d. Confirm the upgrade operation and click **Upgrade**.

      Note

      If a certificate notice is displayed, review and approve it.

   e. Click the **Monitor** tab, and approve each step of the upgrade process.
   f. When complete, click **Mark operation completed**.

After you finish

Complete the upgrade by following the steps described in Upgrading from v2.0.0.2+ with the IM - Windows server on page 43.

Upgrading from v2.0.0.2+ with the IM - Windows server

Use the Installation Manager (IM) to upgrade ScaleIO components on a Windows server, from v2.0.0.2 (or later) to the current version.

Before you begin

Ensure that:

- The EMC ScaleIO gateway service is started.

- The Visual C++ redistributable 2010 package (64-bit) is installed on the Gateway server.
- If the Windows OS is installed on the SATADOM, you must make changes to the memory settings before upgrading, as described in Upgrading Windows servers when the OS is installed on SATADOM on page 176.
- If a server has RFcache installed on it, you must prepare the server BEFORE upgrading. The procedure is described in RFcache is not installed on an upgraded Windows server on page 218.

This procedure can be used for upgrading a 3 or 5-node cluster.

When SNMP is enabled in the 2.x base version, SNMP MDM credentials remain in the v2.x Lockbox after the upgrade.

Procedure

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

   `EMC-ScaleIO-gateway-2.0-13000.X-x64.msi`

2. Run the file.

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

3. Log in to the IM server (https://*<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in "Installing with the Installation Manager".

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

4. Upgrade the ScaleIO components:
   a. From the IM, click the **Maintain** tab.
   b. From the **Maintenance operation** screen, perform the following:
      a. Type the Master (or Primary) MDM credentials.
      b. Type the LIA password.
      c. Click **Retrieve system topology**. A success message is displayed.
      d. Click **Upgrade**.
   c. Set upgrade options:
      - Set performance profiles.
      - Automatic restart of Windows servers.
      - Enable parallel SDS upgrades.
   d. Confirm the upgrade operation and click **Upgrade**.

Note

If a certificate notice is displayed, review and approve it.

e. Click the **Monitor** tab, and approve each step of the upgrade process.

f. When complete, click **Mark operation completed**.

CHAPTER 4

Automated Upgrade of ESXi Servers

This chapter describes how to use the ScaleIO automated management tools to upgrade ScaleIO components on ESXi servers. Topics include:

- Upgrading ESXi configurations 48
- Preparing to upgrade ESXi configurations 48
- Upgrading from ScaleIO v1.32.x - ESXi server 49
- Upgrading from ScaleIO v2.0 or v2.0.0.1 - ESXi server 54
- Upgrading from ScaleIO v2.0.0.2+ - ESXi server 57

Upgrading ESXi configurations

This section describes how to upgrade when all ScaleIO components reside on ESXi servers.

The upgrade uses the Installation Manager (IM, part of the ScaleIO Gateway), together with the LIA of the new version (on an SVM), to orchestrate the upgrade.

Upgrading is non-disruptive; you can upgrade while IOs are running and volumes are mapped.

Extending from a 3-node to a 5-node cluster is performed, and described, after upgrading.

Preparing to upgrade ESXi configurations

When upgrading from v1.32.x, do not delete volumes or snapshots during the MDM cluster upgrade phase of the upgrade. This phase is normally very short.

Before upgrading the ScaleIO components from all previous versions, perform the following:

Procedure

1. Install Java 1.8 on the Gateway and GUI servers. For example:
   - Linux:

     ```
     rpm -i /var/tmp/jre-8uXX-linux-x64.rpm
     ```

   - Windows:

     ```
     jre-8uXX-windows-x64.exe
     ```

2. Download and extract the ScaleIO packages to `/tmp/`.
3. Upgrade the GUI (on the Windows server), by running the following:

   ```
   EMC-ScaleIO-gui-2.0-13000.X.msi
   ```

4. Install the OpenSSL packages on all SVMs:
   a. Copy the SLES 11.3 OpenSSL RPM packages, from the ISO, to each SVM.
   b. Install the files, by running these commands:

   ```
   rpm -i libopenssl1_0_0-1.0.1g-0.40.1.x86_64.rpm
   ```

   ```
   rpm -i openssl1-1.0.1g-0.40.1.x86_64.rpm
   ```

5. When upgrading a server with ESXi 6u2:
   a. Ensure that the following line does not appear in the `etc/ssh/sshd_config` file of the ESX host:

```
kexalgorithms diffie-hellman-group1-sha1,diffie-hellman-
group-exchange-sha1
```

If the referred line appears, then delete it. The line deletion does not relate to any configuration and you do not need to restart either the ESX host or the sshd service.

b. Repeat the procedure for the rest of the servers with ESXi 6u2.

6. The upgrade process requires entering the ESXi servers into maintenance mode. To do so, ensure that Admission Control is not enabled on the VMware cluster:

Client	Procedure
vSphere Web Client	a. From the vSphere Web Client, log in to the vCenter. b. Select the cluster that contains the ScaleIO nodes. c. Select the **Configure** tab. d. Select **Services** > **vSphere Availability**. e. At the bottom of the window, expand **Admission Control**, and ensure that it is **Disabled**.
vSphere Client	a. From the vSphere client, log in to the vCenter. b. Right-click a cluster that contains ScaleIO nodes and select **Edit Settings**. c. Click the **vSphere HA** menu, and ensure that **Admission Control** is set to **Disable**.

After you finish

Continue the upgrade according to your operating system environment.

Upgrading from ScaleIO v1.32.x - ESXi server

This section describes how to upgrade from ScaleIO v1.32.4 (or later). To upgrade from pre-1.32.4 versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.

Unregister the old ScaleIO plug-in

If you unregistered the plug-in already (such as when first upgrading from v.2.0 to v2.0.0.2), you can skip to the next step.

Procedure

1. From the folder where you extracted the `EMC-ScaleIO-vSphere-plugin-installer-1.32-XXX.0.zip` file, use PowerCLI, as administrator, to run the `ScaleIOPluginSetup-1.32-XXX.0.ps1` script.
2. Enter the vCenter credentials and confirm the script actions.
3. Select option **2, Unregister ScaleIO plugin**.
4. Log out, then log back in to the vSphere web client.

   Verify that the plug-in is no longer registered.

5. When the process is complete, enter `4` to exit the plug-in script.

Register the new ScaleIO plug-in

Register the newScaleIO vSphere plug-in.

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.

To use your own web server, see "Manual registration of the ScaleIO plug-in".

Procedure

1. Copy the following files to your designated host (preferably your vCenter):
   - `ScaleIOVM_2nics_2.0.14000.X.ova`
   - `EMC-ScaleIO-vSphere-plugin-installer-2.0-13000.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-2.0-13000.X.ps1`.
   a. Enter the vCenter name or IP address, user name, and password.
   b. For `Choose Mode`, select option **1, Register SIO plugin**.
   c. Read the upgrade notice, and enter `y` to continue.
   d. For `Select Registration Mode`, choose **Standard** (simplified, using embedded Tomcat).

      This step may take a few minutes.
   e. If necessary, accept the thumbprint.
4. Log in to the vSphere web client.
   a. If you are already logged in, log out, then log in again.
   b. Verify that the ScaleIO icon appears in the web client.
   c. Open the plug-in in the client, and ensure that it is the correct version.

      If the version didn't update, clear your browser cookies, log out, then log in to the plug-in again.
5. In the PowerCLI window, press ENTER to finish the plug-in download and return to the menu.

Register the existing system

Procedure

1. From the main plug-in window, click **ScaleIO systems**.
2. Right-click the system to register, and select **Reregister ScaleIO system**.
3. Enter the user name and password of the system, and click **OK**.

   The new plug-in requires a secured connection that was not supported in the previous version. If a secured connection message appears, you must approve non-secured connection to complete reregistration.

Create the datastore templates

Procedure

1. From the vSphere **Home** tab, verify that the ScaleIO icon

is visible in the **Inventories** section.

Note

If the SIO icon is missing, the vSphere web client (Virgo) server failed to download/register the plug-in due to one of the following reasons:

- A connectivity problem (for example, network/firewall etc.) between the vSphere web client server and the web server storing the plug-in. To resolve the issue, verify that there is communication between the vSphere web client server and the web server that is storing the plug-in.
- URL problem. To resolve the issue, verify that the URL is `https://` and is pointing to the correct web server IP address (i.e. ScaleIO Gateway).

For information on how to use the log to troubleshoot problems that may arise, see "Troubleshooting plug-in registration issues".

To upload the OVA template, perform the following:

2. From the PowerCLI script, select **3, Create SVM template**.
3. Enter 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, enter

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 (v2.0.13000.X) was
successfully created
```

4. When the process is complete, enter **4** to exit the plug-in script.

Upgrade the Gateway

Procedure

1. From the vCenter client, shut down the Gateway SVM.
2. Increase the memory on the Gateway SVM to 3072 MB.
3. Power on the Gateway SVM.
4. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

Upgrade the ScaleIO components

Procedure

1. From the plug-in **ScaleIO Systems** screen, launch the Gateway's Installation Manager (IM) by clicking **Open ScaleIO Gateway**.
2. From the IM, enter the default IM credentials, then upload the following packages, as described in "Upload installation packages":

Note

Pay close attention to the file names, as they are from different versions, and in different formats.

Packages from base version:

- `EMC-ScaleIO-lia-1.32-XXX.X.<flavor>.SVM.x86_64.rpm`
- `EMC-ScaleIO-mdm-1.32-XXX.X.<flavor>.SVM.x86_64.rpm`
- `EMC-ScaleIO-sdc-1.32-XXX.X.<flavor>.SVM.x86_64.rpm`
- `EMC-ScaleIO-sds-1.32-XXX.X.<flavor>.SVM.x86_64.rpm`
- `EMC-ScaleIO-tb-1.32-XXX.X.<flavor>.SVM.x86_64.rpm`
- `EMC-ScaleIO-callhome-1.32-XXX.X.<flavor>.SVM.x86_64.rpm`

Packages from new version:

- `EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm`
- `EMC-ScaleIO-sds-2.0-13000.X.<flavor>.x86_64.rpm`
- `EMC-ScaleIO-lia-2.0-13000.X.<flavor>.x86_64.rpm`

3. Upgrade the ScaleIO components:
   a. From the IM, click the **Maintain** tab.
   b. From the **Maintenance operation** screen, perform the upgrade:
      a. Type the Master (or Primary) MDM credentials.
      b. Type the LIA password.
      c. From the **Advanced Options** section, select both of the **Allow non-secure communication** options.

         Security warnings are displayed because the new Gateway is trying to communicate in a non-secure mode. You can ignore the messages.
      d. Click **Retrieve system topology**. A success message is displayed.
      e. Click **Upgrade**.
   c. Set upgrade options:
      - Set performance profiles.
      - Automatic restart of Windows servers.
      - Enable parallel SDS upgrades.
   d. Confirm the upgrade operation and click **Upgrade**.

      Note

      If a certificate notice is displayed, review and approve it. You may need to rerun the phase to continue.

   e. Click the **Monitor** tab, and approve each step of the upgrade process.
   f. When complete, click **Mark operation completed**.

Upgrade and restart the SDCs

This process involves restarting the SDC server, which may cause IO failures. It is recommended to migrate VMs before starting.

Procedure

1. From the VMware plug-in main screen, click **SDCs**.
2. On the **SDCs** screen, right-click an SDC and select **Upgrade SDC**.

3. In the **Upgrade SDC** dialog box, perform the following for each ESX:
   a. Enter the ESX password and click **OK**.
   b. When the upgrade completes, click **Finish**, then **Close**.
4. From the vCenter, reboot each ESX, one at a time.

   After each ESX reboot, ensure that there is no rebuild or rebalance in process before rebooting the next server.

   Note

   After upgrading from 1.32.x, you must perform manual memory allocation on the SVM, as described in "SVM manual memory allocation".

After upgrading

These upgrade procedures leave the system in non-secure authentication mode. After the upgrade, you can switch to secure mode, as well as perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node

Note

If you choose not to switch to secure mode, you need to disable secure communication on every MDM server before you can run SCLI commands.

These options, and others, are described in "Maintaining a ScaleIO system".

Upgrading from ScaleIO v2.0 or v2.0.0.1 - ESXi server

This section describes how to upgrade from ScaleIO v2.0 or v2.0.0.1 to the current version. This procedure can be used to upgrade from 3-node or 5-node cluster mode.

This procedure is completed with the following steps:

1. Upgrade to v2.0.0.x (latest) - described here.
2. Upgrade from v2.0.0.x (latest) to the current version - described in Upgrading from ScaleIO v2.0.0.2+ - ESXi server on page 57.

Unregister the old ScaleIO plug-in

If you unregistered the plug-in already (such as when first upgrading from v.2.0 to v2.0.0.x (latest)), you can skip to the next step.

Procedure

1. From the folder where you extracted the `EMC-ScaleIO-vSphere-plugin-installer-2.0-XXX.X.zip` file, use PowerCLI, as administrator, to run the `ScaleIOPluginSetup-2.0-XXX.X.ps1` script.
2. Enter the vCenter credentials and confirm the script actions.
3. Select option **2, Unregister ScaleIO plugin**.
4. Log out, then log back in to the vSphere web client.

   Verify that the plug-in is no longer registered.

5. When the process is complete, enter `4` to exit the plug-in script.

Register the new ScaleIO plug-in

Register the ScaleIO vSphere plug-in.

Note

To use your own web server, see "Manual registration of the ScaleIO plug-in".

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.

Procedure

1. Copy the following files to your designated host (preferably your vCenter):
   - `ScaleIOVM_2nics_2.0.XXX.X.ova`
   - `EMC-ScaleIO-vSphere-plugin-installer-2.0-XXX.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-2.0-XXX.X.ps1`.
   a. Enter the vCenter name or IP address, user name, and password.
   b. For `Choose Mode`, select option **1, Register SIO plugin**.
   c. Read the upgrade notice, and enter `y` to continue.
   d. For `Select Registration Mode`, choose **Standard** (simplified, using embedded Tomcat).

      This step may take a few minutes.
   e. If necessary, accept the thumbprint.
4. Log in to the vSphere web client.
   a. If you are already logged in, log out, then log in again.
   b. Verify that the ScaleIO icon appears in the web client.
   c. Open the plug-in in the client, and ensure that it is the correct version.

      If the version didn't update, clear your browser cookies, log out, then log in to the plug-in again.
5. In the PowerCLI window, press ENTER to finish the plug-in download and return to the menu.

Register the existing system

Procedure

1. From the main plug-in window, click **ScaleIO systems**.
2. Right-click the system to register, and select **Reregister ScaleIO system**.
3. Enter the user name and password of the system, and click **OK**.

Upgrade the Gateway

Procedure

1. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-gateway-2.0-
XXX.0.x86_64.rpm (for SLES 11.3 add --nodeps)
```

Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-XXX.0.x86_64.rpm (--nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

Upgrade the ScaleIO components

Procedure

1. From the plug-in **ScaleIO Systems** screen, launch the Gateway's Installation Manager (IM) by clicking **Open ScaleIO Gateway**.
2. From the IM, enter the default IM credentials, then upload the following packages, as described in "Upload installation packages":

   Packages from base version: lia, mdm, sdc, sds, xcache (optional), for example:

   - `EMC-ScaleIO-lia-2.0-XXX.0.<flavor>.x86_64.rpm`
   - `EMC-ScaleIO-mdm-2.0-XXX.0.<flavor>.x86_64.rpm`

   Packages from new version: lia, mdm, sdc, sds, xcache (optional), for example:

   - `EMC-ScaleIO-lia-2.0-XXX.0.<flavor>.x86_64.rpm`
   - `EMC-ScaleIO-mdm-2.0-XXX.0.<flavor>.x86_64.rpm`

3. Upgrade the ScaleIO components:
   a. From the IM, click the **Maintain** tab.
   b. From the **Maintenance operation** screen, perform the following:
      a. Type the Master (or Primary) MDM credentials.
      b. Type the LIA password.
      c. Click **Retrieve system topology**. A success message is displayed.
      d. Click **Upgrade**.
   c. Set upgrade options:
      - Set performance profiles.
      - Automatic restart of Windows servers.
      - Enable parallel SDS upgrades.
   d. Confirm the upgrade operation and click **Upgrade**.

Note

If a certificate notice is displayed, review and approve it.

e. Click the **Monitor** tab, and approve each step of the upgrade process.

f. When complete, click **Mark operation completed**.

After you finish

To continue the upgrade to v.2.0.1, perform the procedures described in Upgrading from ScaleIO v2.0.0.2+ - ESXi server on page 57.

Upgrading from ScaleIO v2.0.0.2+ - ESXi server

This section describes how to upgrade from ScaleIO v2.0.0.2 (or later). This procedure can be used to upgrade from 3-node or 5-node cluster mode.

Unregister the old ScaleIO plug-in

If you unregistered the plug-in already (such as when first upgrading from v.2.0 to v2.0.0.x), you can skip to the next step.

Procedure

1. From the folder where you extracted the `EMC-ScaleIO-vSphere-plugin-installer-2.0-XXXX.X.zip` file, use PowerCLI, as administrator, to run the `ScaleIOPluginSetup-2.0-XXXX.X.ps1` script.
2. Enter the vCenter credentials and confirm the script actions.
3. Select option **2, Unregister ScaleIO plugin**.
4. Log out, then log back in to the vSphere web client.

   Verify that the plug-in is no longer registered.
5. When the process is complete, enter `4` to exit the plug-in script.

Register the new ScaleIO plug-in

Register the ScaleIO v2.0.1 vSphere plug-in.

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.

To use your own web server, see "Manual registration of the ScaleIO plug-in".

Procedure

1. Copy the following files to your designated host (preferably your vCenter):
   - `ScaleIOVM_2nics_2.0.13000.X.ova`
   - `EMC-ScaleIO-vSphere-plugin-installer-2.0-13000.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-2.0-13000.X.ps1`.

a. Enter the vCenter name or IP address, user name, and password.

b. For `Choose Mode`, select option **1, Register SIO plugin**.

c. Read the upgrade notice, and enter `y` to continue.

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

This step may take a few minutes.

e. If necessary, accept the thumbprint.

4. Log in to the vSphere web client.

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

b. Verify that the ScaleIO icon appears in the web client.

c. Open the plug-in in the client, and ensure that it is the correct version.

If the version didn't update, clear your browser cookies, log out, then log in to the plug-in again.

5. In the PowerCLI window, press **Enter** to finish the plug-in download and return to the menu.

Register the existing system

Procedure

1. From the main plug-in window, click **ScaleIO systems**.
2. Right-click the system to register, and select **Reregister ScaleIO system**.
3. Enter the user name and password of the system, and click **OK**.

Create the datastore templates

Procedure

1. From the vSphere **Home** tab, verify that the ScaleIO icon

is visible in the **Inventories** section.

Note

If the SIO icon is missing, the vSphere web client (Virgo) server failed to download/register the plug-in due to one of the following reasons:

- A connectivity problem (for example, network/firewall etc.) between the vSphere web client server and the web server storing the plug-in. To resolve the issue, verify that there is communication between the vSphere web client server and the web server that is storing the plug-in.
- URL problem. To resolve the issue, verify that the URL is `https://` and is pointing to the correct web server IP address (i.e. ScaleIO Gateway).

For information on how to use the log to troubleshoot problems that may arise, see "Troubleshooting plug-in registration issues".

To upload the OVA template, perform the following:

2. From the PowerCLI script, select **3, Create SVM template**.
3. Enter 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, enter 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 (v2.0.13000.X) was
successfully created
```

4. When the process is complete, enter `4` to exit the plug-in script.

Upgrade the Gateway

Before you begin

Ensure that you have the credentials to access the Gateway server.

Procedure

1. Use scp to copy the Gateway package to `/tmp`.
2. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

Upgrade the ScaleIO components

Procedure

1. From the plug-in **EMC ScaleIO Systems** screen, launch the Gateway's Installation Manager (IM) by clicking **Open ScaleIO Gateway**.
2. From the IM, enter the default IM credentials, then upload the following packages, as described in "Upload installation packages":

   Packages from base version: lia, mdm, sds, xcache (optional), for example:

   - `EMC-ScaleIO-lia-2.0-XXX.0.<flavor>.x86_64.rpm`
   - `EMC-ScaleIO-mdm-2.0-XXX.0.<flavor>.x86_64.rpm`

   Packages from new version: lia, mdm, sds, xcache (optional), for example:

   - `EMC-ScaleIO-lia-2.0-13000.X.<flavor>.x86_64.rpm`
   - `EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm`

3. Upgrade the ScaleIO components:

   a. From the IM, click the **Maintain** tab.

   b. From the **Maintenance operation** screen, perform the following:

      a. Type the Master (or Primary) MDM credentials.

      b. Type the LIA password.

      c. Click **Retrieve system topology**. A success message is displayed.

      d. Click **Upgrade**.

   c. Set upgrade options:

      - Set performance profiles.
      - Enable parallel SDS upgrades.

   d. Confirm the upgrade operation and click **Upgrade**.

   Note

   If a certificate notice is displayed, review and approve it.

   e. Click the **Monitor** tab, and approve each step of the upgrade process.

   f. When complete, click **Mark operation completed**.

After you finish

After upgrade, you can add virtual IP addresses. For more information, see Configuring virtual IP addresses on page 196.

Upgrade and restart the SDCs

This process involves restarting the SDC server, which may cause IO failures. It is recommended to migrate VMs before starting.

Procedure

1. From the VMware plug-in main screen, click **SDCs**.
2. On the **SDCs** screen, right-click an SDC and select **Upgrade SDC**.
3. In the **Upgrade SDC** dialog box, perform the following for each ESX:
   a. Enter the ESX password and click **OK**.
   b. When the upgrade completes, click **Finish**, then **Close**.
4. From the vCenter, reboot each ESX, one at a time.

   After each ESX reboot, ensure that there is no rebuild or rebalance in process before rebooting the next server.

CHAPTER 5

Automated Upgrade of Servers in a 2-Layer Configuration

This chapter describes how to use the ScaleIO automated management tools to upgrade ScaleIO components in a 2-layer configuration. Topics include:

- Preparing to upgrade 2-layer configuration - Linux 64
- Upgrading 2-layer configurations (Linux and ESX) 64

Preparing to upgrade 2-layer configuration - Linux

Before upgrading the ScaleIO components, perform the following steps:

Procedure

1. Install Java 1.8 on the Gateway and GUI servers. For example:
   - Linux:

     ```
     rpm -i /var/tmp/jre-8uXX-linux-x64.rpm
     ```

   - Windows:

     ```
     jre-8uXX-windows-x64.exe
     ```

2. Download and extract the ScaleIO packages to `/tmp/`
3. Upgrade the GUI:

   ```
   EMC-ScaleIO-gui-2.0-13000.X.msi
   ```

After you finish

The next upgrade steps for your environment are determined by the base version of ScaleIO and whether the LIA is installed. Proceed to the section that matches your environment.

Upgrading 2-layer configurations (Linux and ESX)

This section describes how to upgrade when the MDM and SDS reside on bare metal servers, and the SDCs reside on ESXi servers.

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.

ESXi upgrade uses the VMware plug-in.

Note

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 already installed, you will first upgrade it to the new version. If no LIA is installed, you must install the new version.

Upgrading is non-disruptive; you can upgrade while IOs are running and volumes are mapped.

Extending from a 3-node to a 5-node cluster is performed, and described, after upgrading.

Upgrading v1.32.x 2-layer (Linux) with LIA installed

This section describes how to upgrade from ScaleIO v1.32.4 (or later), when LIA is installed, on servers in a Linux 2-layer environment. To upgrade from pre-1.32.4

versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.

Procedure

1. Edit the following file:

   `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/gatewayUser.properties`

2. Add the complete list of SDC IP addresses to `im.ip.ignore.list`.
3. Restart the Gateway, by running this command:

   ```
   /etc/init.d/scaleio-gateway restart
   ```

4. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

   ```
   SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
   gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
   ```

   Example:

   ```
   SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
   ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
   ```

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

5. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

   ```
   SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
   gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
   ```

   Example:

   ```
   SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
   ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
   ```

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

6. Log in to the IM server (https://*<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in "Installing with the Installation Manager".

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

7. Upgrade the ScaleIO components:

a. From the IM, click the **Maintain** tab.

b. From the **Maintenance operation** screen, perform the upgrade:

   a. Type the Master (or Primary) MDM credentials.

   b. Type the LIA password.

   c. From the **Advanced Options** section, select both of the **Allow non-secure communication** options.

      Security warnings are displayed because the new Gateway is trying to communicate in a non-secure mode. You can ignore the messages.

   d. Click **Retrieve system topology**. A success message is displayed.

   e. Click **Upgrade**.

c. Set upgrade options:

   - Set performance profiles.
   - Automatic restart of Windows servers.
   - Enable parallel SDS upgrades.

d. Confirm the upgrade operation and click **Upgrade**.

Note

If a certificate notice is displayed, review and approve it. You may need to rerun the phase to continue.

---

e. Click the **Monitor** tab, and approve each step of the upgrade process.

f. When complete, click **Mark operation completed**.

8. Remove the old plug-in:

   a. the folder where you extracted the `EMC-ScaleIO-vSphere-plugin-installer-1.32-XXX.X.zip` file, use PowerCLI to run the `ScaleIOPluginSetup-1.32-XXX.X.ps1 script`.

   b. Select option **2, Unregister ScaleIO plugin**.

   c. Enter the vCenter credentials and confirm the script actions.

   d. Log out, then log back in to the vSphere web client.

      The plug-in is no longer registered.

9. Register the new plug-in, as described in Register the new ScaleIO plug-in on page 55.

10. Upgrade SDCs on the ESXi servers, as described in Upgrade and restart the SDCs on page 53.

    This step requires restarting the SDC server, which may cause IO failures. It is recommended to migrate VMs before starting.

2-layer (Linux)—v1.32.x, no LIA installed

This section describes how to upgrade from ScaleIO, from v1.32.4 (or later), when LIA is not installed. To upgrade from pre-1.32.4 versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.

The first step is to install the LIA; then proceed to upgrade the rest of the system.

Procedure

1. Edit the following file:

   `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/gatewayUser.properties`

2. Add the complete list of SDC IP addresses to `im.ip.ignore.list`.
3. Restart the Gateway, by running this command:

   ```
   /etc/init.d/scaleio-gateway restart
   ```

4. 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-2.0-13000.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 (! @#$ ...)

5. 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.

6. Restart the LIA service by running the following command:

   ```
   pkill lia
   ```

7. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

   ```
   SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
   gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
   ```

Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

8. Log in to the IM server (https://<IM_Server_IP>), and for each OS in your ScaleIO environment, upload the installation packages, as described in "Installing with the Installation Manager".

Note

Ensure that the v1.32.x MDM and TB packages are present. If they are not, upload them, also.

---

9. Upgrade the ScaleIO components:
   a. From the IM, click the **Maintain** tab.
   b. From the **Maintenance operation** screen, perform the upgrade:
      a. Type the Master (or Primary) MDM credentials.
      b. Type the LIA password.
      c. From the **Advanced Options** section, select both of the **Allow non-secure communication** options.

         Security warnings are displayed because the new Gateway is trying to communicate in a non-secure mode. You can ignore the messages.
      d. Click **Retrieve system topology**. A success message is displayed.
      e. Click **Upgrade**.
   c. Set upgrade options:
      - Set performance profiles.
      - Automatic restart of Windows servers.
      - Enable parallel SDS upgrades.
   d. Confirm the upgrade operation and click **Upgrade**.

      Note

      If a certificate notice is displayed, review and approve it. You may need to rerun the phase to continue.

      ---
   e. Click the **Monitor** tab, and approve each step of the upgrade process.
   f. When complete, click **Mark operation completed**.
10. The following steps involve restarting the SDC server, which may cause IO failures. It is recommended to migrate VMs before starting.
11. Remove the old plug-in:

a. the folder where you extracted the `EMC-ScaleIO-vSphere-plugin-installer-1.32-XXX.X.zip` file, use PowerCLI to run the `ScaleIOPluginSetup-1.32-XXX.X.ps1 script`.

b. Select option **2, Unregister ScaleIO plugin**.

c. Enter the vCenter credentials and confirm the script actions.

d. Log out, then log back in to the vSphere web client.

The plug-in is no longer registered.

12. Register the new plug-in, as described in Register the new ScaleIO plug-in on page 55.
13. Upgrade SDCs on the ESXi servers, as described in Upgrade and restart the SDCs on page 53.

    This step requires restarting the SDC server, which may cause IO failures. It is recommended to migrate VMs before starting.

Upgrading v2.0 or v2.0.0.1 2-layer

Upgrade from ScaleIO v2.0 or v2.0.0.1 to the current version, in a 2-layer Linux environment.

This procedure can be used to upgrade a 3-node or 5-node cluster.

This procedure is completed with the following steps:

1. Upgrade from the base version to v2.0.0.x (latest) - described here.
2. Upgrade from v2.0.0.x (latest) to the current version - described in Upgrading v2.0.0.2+ 2-layer - Linux server on page 71.
   Begin the upgrade:

Procedure

1. Edit the following file:

   `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/gatewayUser.properties`

2. Add the complete list of SDC IP addresses to `im.ip.ignore.list`.
3. Restart the Gateway:

```
/etc/init.d/scaleio-gateway restart
```

4. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

   Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

5. Log in to the IM server (https://*<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in "Installing with the Installation Manager".

Note

To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

---

6. Upgrade the ScaleIO components:
   a. From the IM, click the **Maintain** tab.
   b. From the **Maintenance operation** screen, perform the following:
      a. Type the Master (or Primary) MDM credentials.
      b. Type the LIA password.
      c. Click **Retrieve system topology**. A success message is displayed.
      d. Click **Upgrade**.
   c. Set upgrade options:
      - Set performance profiles.
      - Automatic restart of Windows servers.
      - Enable parallel SDS upgrades.
   d. Confirm the upgrade operation and click **Upgrade**.

      Note

      If a certificate notice is displayed, review and approve it.

      ---

   e. Click the **Monitor** tab, and approve each step of the upgrade process.
   f. When complete, click **Mark operation completed**.
7. Remove the old plug-in:
   a. From the folder where you extracted the `EMC-ScaleIO-vSphere-plugin-installer-2.0-X.X.zip` file, use PowerCLI, as administrator, to run the `ScaleIOPluginSetup-2.0-XXX.X.ps1` script.
   b. Enter the vCenter credentials and confirm the script actions.
   c. Select option **2, Unregister ScaleIO plugin**.
   d. Log out, then log back in to the vSphere web client.

      Verify that the plug-in is no longer registered.
   e. When the process is complete, enter `4` to exit the plug-in script.
8. Register the new plug-in, as described in Register the new ScaleIO plug-in on page 55.
9. Upgrade SDCs on the ESXi servers, as described in Upgrade and restart the SDCs on page 53.

This step requires restarting the SDC server, which may cause IO failures. It is recommended to migrate VMs before starting.

After you finish

You can add virtual IP addresses as described in Configuring virtual IP addresses on page 196.

Upgrading v2.0.0.2+ 2-layer - Linux server

Upgrade from ScaleIO v2.0.0.2 (or later) to the current version, in a Linux environment.

This procedure can be used to upgrade a 3-node or 5-node cluster.

Procedure

1. Edit the following file:

   `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/gatewayUser.properties`

2. Add the complete list of SDC IP addresses to `im.ip.ignore.list`.
3. Restart the Gateway, by running this command:

   ```
   /etc/init.d/scaleio-gateway restart
   ```

4. Upgrade the currently installed (or install a new) Linux Gateway (run command all on one line):

   ```
   SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
   gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
   ```

   Example:

   ```
   SIO_GW_KEYTOOL=/usr/java/jre1.8.0_XX/bin/ rpm -U /tmp/EMC-
   ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
   ```

   A success message is displayed, and the IM interface now reflects the updated Gateway version.

5. Log in to the IM server (https://*<IM_Server_IP>*), and for each OS in your ScaleIO environment, perform the "Upload installation packages" steps (only), as described in "Installing with the Installation Manager".

   Note

   To enable roll-back, your current packages must be on the IM. You may need to upload them for this option to work.

6. Upgrade the ScaleIO components:
   a. From the IM, click the **Maintain** tab.
   b. From the **Maintenance operation** screen, perform the following:
      a. Type the Master (or Primary) MDM credentials.

b. Type the LIA password.

c. Click **Retrieve system topology**. A success message is displayed.

d. Click **Upgrade**.

c. Set upgrade options:

- Set performance profiles.
- Automatic restart of Windows servers.
- Enable parallel SDS upgrades.

d. Confirm the upgrade operation and click **Upgrade**.

Note

If a certificate notice is displayed, review and approve it.

---

e. Click the **Monitor** tab, and approve each step of the upgrade process.

f. When complete, click **Mark operation completed**.

7. The following steps involve restarting the SDC server, which may cause IO failures. It is recommended to migrate VMs before starting.
8. Remove the old plug-in:

   a. the folder where you extracted the `EMC-ScaleIO-vSphere-plugin-installer-2.0-XXX.X.zip` file, use PowerCLI to run the `ScaleIOPluginSetup-2.0-XXX.X.ps1 script`.

   b. Select option **2, Unregister ScaleIO plugin**.

   c. Enter the vCenter credentials and confirm the script actions.

   d. Log out, then log back in to the vSphere web client.

   The plug-in is no longer registered.

9. Register the new plug-in, as described in Register the new ScaleIO plug-in on page 55.
10. Upgrade SDCs on the ESXi servers, as described in Upgrade and restart the SDCs on page 53.

    This step requires restarting the SDC server, which may cause IO failures. It is recommended to migrate VMs before starting.

PART 3

Manual Upgrade

This part describes how to upgrade ScaleIO components manually. Chapters include:

Chapter 6, "Manual Upgrade of Linux Server"

Chapter 7, "Manual Upgrade of Windows Server"

Chapter 8, "Manual Upgrade of Xen Server"

CHAPTER 6

Manual Upgrade of Linux Server

This chapter describes how to perform a manual upgrade of ScaleIO components on a Linux server. Topics include:

- Manual upgrade of physical Linux configurations 76
- Manual upgrade of Linux system from v1.32.x 77
- Manual upgrade of Linux system from v2.0 or v2.0.0.1, in a 3-node cluster 83
- Manual upgrade of Linux system from v2.0 or v2.0.0.1, in a 5-node cluster 83
- Manual upgrade of Linux system from v2.0.0.2+, in a 3-node cluster 84
- Manual upgrade of Linux systems from v2.0.0.2+, in a 5-node cluster 89

Manual upgrade of physical Linux configurations

This section describes how to perform a manual upgrade when all ScaleIO components reside on physical Linux servers.

It is highly recommended to perform upgrades using the Installation Manager.

Note

When performing a manual upgrade, it is crucial to follow the described upgrade procedures step by step. Failure to do so might result in data loss and system instability.

Upgrading is non-disruptive; you can upgrade while IOs are running and volumes are mapped.

Proceed to the section that matches your system environment.

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>-2.0-13000.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 4 Linux package formats

Linux flavor	Package format	Example
CoreOS	CoreOS	`EMC-ScaleIO-mdm-2.0-13000.X.CoreOS.x86_64.tar` CoreOS packages may need to be extracted before use. This is described where relevant.
RHEL/OL/CentOS	el<version>	`EMC-ScaleIO-mdm-2.0-13000.X.el<version>.x86_64.rpm` Example: `EMC-ScaleIO-mdm-2.0-13000.X.el6.x86_64.rpm`
SUSE	sles<version>	`EMC-ScaleIO-mdm-2.0-13000.X.sles<version>.x86_64.rpm` Example: `EMC-ScaleIO-mdm-2.0-13000.X.sles11.3.x86_64.rpm`
Ubuntu	Ubuntu.<version>	`EMC-ScaleIO-mdm-2.0-13000.X.Ubuntu.<version>.x86_64.tar` Example: `EMC-ScaleIO-mdm-2.0-13000.X.Ubuntu.16.04.x86_64.tar` Ubuntu packages may need to be extracted before use. This is described where relevant.
XenServer	xs<version>	`EMC-ScaleIO-mdm-2.0-13000.X.xs<version>.x86_64.rpm` Example: `EMC-ScaleIO-mdm-2.0-13000.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 upgrading CoreOs, Oracle Linux (OL), or Ubuntu systems, ensure that the ScaleIO environment is prepared, as described in Deploying on Ubuntu, Oracle Linux, or CoreOS servers on page 180.

Manual upgrade of Linux system from v1.32.x

Manual upgrade of ScaleIO v1.32.4 (or later) components, on a Linux server. To upgrade from pre-1.32.4 versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.

Before you begin

- Ensure that there is at least 1.1 GB of available disk space on all MDMs (additional 0.1 GB is required if the server also has an SDS).
- Java 1.8 is required on the Gateway and the GUI servers before performing the upgrade. Download the latest JRE for your operating system from the Oracle website.
- If needed, install the v2.0.1 `GPG-RPM-KEY` file (with the `rpm --import <path_to_new_public_key>` command). This is used to authenticate the RPM packages.
- To upgrade Ubuntu or CoreOS servers, you must first extract all the packages, as described in Extracting ScaleIO packages on page 160.

This procedure can be used for all Linux flavor servers, with the exception of Xen servers. For Xen servers, follow the procedure described in Manual upgrade of Xen system from v1.32.x on page 112.

When SNMP is enabled in the 1.32.x base version, the previous SNMP MDM credentials are erased from the old `gatewayUser.properties` configuration file and moved into the v2.x Lockbox. To harden the password, you can use the SioGWTool after the upgrade, or you can pass an environment variable while upgrading the Gateway. For more information, see "Non-Disruptive Upgrade Issues" in the *ScaleIO Upgrade Guide*.

When upgrading from v1.32.x, do not delete volumes or snapshots during the MDM cluster upgrade phase of the upgrade. This phase is normally very short.

For the purposes of this procedure, we will use the following nomenclature:

MDM 1—Primary MDM

MDM 2—Secondary MDM

TB 1—Tie Breaker

Procedure

1. Download and extract the ScaleIO packages to `/tmp/`.
2. Upgrade the GUI (on the Windows server):

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

3. Remove the v1.32 CallHome component, by running the following command on MDM 1 and MDM 2:

This step does not need to be done for upgrades on CoreOS servers.

```
rpm -e `rpm -qa|grep ScaleIO|grep callhome`
```

4. Upgrade (or install) LIA on all machines in the system:

```
rpm -U EMC-ScaleIO-lia-2.0-13000.X.<flavor>.x86_64.rpm
```

5. Switch to MDM single mode, by running the following commands on MDM 1:

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

```
 scli --switch_to_single_mode
```

6. Remove the Secondary MDM and the Tie Breaker, by running the following commands on MDM 1:

```
scli --remove_secondary_mdm
```

```
scli --remove_tb
```

7. Upgrade MDM 2, by running the following commands on it:
   a. Stop service:

```
/opt/emc/scaleio/mdm/bin/delete_service.sh
```

   b. Set role:

```
echo actor_role_is_manager=1 >> /opt/emc/scaleio/mdm/cfg/
conf.txt
```

   c. Upgrade RPM:

```
rpm -U EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm
```

   d. Start service:

```
/opt/emc/scaleio/mdm/bin/create_service.sh
```

8. Upgrade the Tie Breaker, by running the following commands on TB 1:
   a. Stop service:

```
/opt/emc/scaleio/tb/bin/delete_service.sh
```

b. Uninstall previous version:

```
rpm -e `rpm -qa|grep ScaleIO|grep tb`
```

c. Install new version (with MDM, not TB package):

```
rpm -i EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm
```

9. Add the upgraded MDM 2 and TB 1, by running the following commands on MDM 1:

   a. Add the Secondary MDM:

```
scli --add_secondary_mdm --secondary_mdm_ip <MDM2 IP>
```

   b. Add the TB:

```
scli --add_tb --tb_ip <TB1 IP>
```

10. Switch to cluster mode, by running the following commands on MDM 1:

```
scli --switch_to_cluster_mode
```

Wait for the repository sync to complete, indicated by its state being normal.

```
scli --query_cluster
```

11. Switch MDM ownership, by running the following command on MDM 1:

```
scli --switch_mdm_ownership
```

12. Modify the new Master MDM (MDM2), by running the following commands on MDM 2:

   a. Log in:

```
scli --login --username <USER> --password <PASSWORD> --
use_nonsecure_communication
```

   b. Enable commands:

```
scli --allow_commands_during_upgrade --
use_nonsecure_communication
```

c. View cluster information:

```
scli --query_cluster --use_nonsecure_communication
```

d. Switch to single mode:

```
scli --switch_cluster_mode --cluster_mode 1_node --
remove_slave_mdm_ip <MDM1 IP> --remove_tb_ip <TB1 IP> --
use_nonsecure_communication
```

e. Remove the standby MDM:

```
scli --remove_standby_mdm --remove_mdm_ip <MDM1 IP> --
use_nonsecure_communication
```

f. Modify the management IP address (run the command on one line):

```
scli --modify_management_ip --new_mdm_management_ip <MDM2
management IP> --target_mdm_ip <MDM2 IP>
--allow_duplicate_management_ips --
use_nonsecure_communication
```

13. Upgrade the MDM 1, by running the following commands on MDM 1:

a. Stop service:

```
/opt/emc/scaleio/mdm/bin/delete_service.sh
```

b. Set role:

```
echo actor_role_is_manager=1 >> /opt/emc/scaleio/mdm/cfg/
conf.txt
```

c. Upgrade RPM:

```
rpm -U EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm
```

d. Start service:

```
/opt/emc/scaleio/mdm/bin/create_service.sh
```

14. Add the MDM as a standby MDM, by running these commands from the Master MDM (MDM 2):

a. Log in (if not logged in):

```
scli --login --username <USER> --password <PASSWORD> --
use_nonsecure_communication
```

b. Enable commands (if not already done):

```
scli --allow_commands_during_upgrade --
use_nonsecure_communication
```

c. Add the standby MDM:

```
scli --add_standby_mdm --new_mdm_ip <MDM1 IP> --mdm_role
manager --new_mdm_management_ip <MDM1 management IP> --
new_mdm_name <MDM1 name> --use_nonsecure_communication
```

15. Switch to 3-node cluster mode, by running the following commands, from MDM 2:

a. Switch mode:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_ip <MDM1 IP> --add_tb_ip <TB1 IP> --
use_nonsecure_communication
```

b. Wait for the cluster to finish synching the repository:

```
scli --query_cluster --use_nonsecure_communication
```

Wait until cluster state is Normal.

c. Switch cluster ownership:

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM1 IP> --
use_nonsecure_communication
```

16. Upgrade all SDSs, by running the following command on each SDS server:

```
rpm -U EMC-ScaleIO-sds-2.0-13000.X.<flavor>.x86_64.rpm
```

You can update (and reboot) all SDSs in the same Fault Set. You can upgrade (and reboot) all SDSs or Fault Sets in the same Protection Domain.

Note

Between the upgrade of each set of SDSs, wait until all rebuilds are complete, and five seconds longer.

17. Upgrade all SDCs, by running the following commands on each SDC server:

```
rpm -U EMC-ScaleIO-sdc-2.0-13000.X.<flavor>.x86_64.rpm
```

If the SDC kernel module is processing IO, its upgrade will occur only after the server is restarted.

18. Validate the upgrade, by running the following command on MDM 1:

```
scli --query_upgrade --use_nonsecure_communication
```

19. Finalize the upgrade, by running the following command on MDM 1:

```
scli --finalize_upgrade --use_nonsecure_communication
```

20. Upgrade the currently installed Gateway by running the following command (optional):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_40/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

NOTICE

If your gateway uses a non-default certificate, before upgrading the gateway you must copy the certificate and other files. For more information, see "Upgrading the Gateway when a non-default certificate is used".

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 148.

After upgrading

These upgrade procedures leave the system in non-secure authentication mode. After the upgrade, you can switch to secure mode, as well as perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses

These options, and others, are described in "Maintaining a ScaleIO system".

Note

If you choose not to switch to secure mode, you need to disable secure communication on every MDM server before you can run SCLI commands.

Manual upgrade of Linux system from v2.0 or v2.0.0.1, in a 3-node cluster

Manual upgrade of ScaleIO v2.0 or v2.0.0.1 components, on a Linux server (all flavors) when the MDM cluster is in 3-node mode.

This procedure is completed as follows:

- First, upgrade to v2.0.0.x (latest), using the same steps described in Manual upgrade of Linux system from v2.0.0.2+, in a 3-node cluster on page 84.
  In the upload packages step, upload the base version and the v2.0.0.x (latest) packages.
- Then, upgrade from v2.0.0.x (latest) to the current version, using those exact same steps.
  In the upload packages step, upload the v2.0.0.x (latest) and the current version packages.

After upgrading

After the upgrade, you can perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses

These options, and others, are described in "Maintaining a ScaleIO system".

Manual upgrade of Linux system from v2.0 or v2.0.0.1, in a 5-node cluster

Manual upgrade of ScaleIO v2.0 or v2.0.0.1 components, on a Linux server (all flavors) when the MDM cluster is in 5-node mode.

This procedure is completed as follows:

- First, upgrade to v2.0.0.x (latest), using the same steps described in Manual upgrade of Linux systems from v2.0.0.2+, in a 5-node cluster on page 89.
  In the upload packages step, upload the base version and the v2.0.0.x (latest) packages.
- Then, upgrade from v2.0.0.x (latest) to the current version, using those exact same steps.
  In the upload packages step, upload the v2.0.0.x (latest) and the current version packages.

After upgrading

After the upgrade, you can perform many other tasks, including:

- Install RFcache

- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses

These options, and others, are described in "Maintaining a ScaleIO system".

Manual upgrade of Linux system from v2.0.0.2+, in a 3-node cluster

Manual upgrade of ScaleIO v2.0.0.2 (or later) components, on a Linux server (all flavors) when the MDM cluster is in 3-node mode.

Before you begin

- Ensure that there is at least 1.1 GB of available disk space on all MDMs (additional 0.1 GB is required if the server also has an SDS).
- Java 1.8 is required on the Gateway and the GUI servers before performing the upgrade. Download the latest JRE for your operating system from the Oracle website.
- If needed, install the v2.0.1 `GPG-RPM-KEY` file (with the `rpm --import <path_to_new_public_key>` command). This is used to used to authenticate the RPM packages.
- To upgrade Ubuntu or CoreOS servers, you must first extract all the packages, as described in Extracting ScaleIO packages on page 160.
- Run the ScaleIO system analysis to analyze the ScaleIO system immediately prior to upgrading. It will identify and avoid faulty setups in the system. For more information, see System analysis overview on page 148.

When SNMP is enabled in the 2.x base version, SNMP MDM credentials remain in the v2.x Lockbox after the upgrade.

When performing this procedure on Xen server, use `siocli` in place of `scli`.

For the purposes of this procedure, we will use the following nomenclature:

MDM 1—Primary MDM

MDM 2—Secondary MDM

TB 1—Tie Breaker

Procedure

1. Download and extract the ScaleIO packages to `/tmp/`.
2. Upgrade the GUI (on the Windows server):

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

3. Upgrade (or install) LIA on all machines in the system:

```
rpm -U EMC-ScaleIO-lia-2.0-13000.X.elX.x86_64.rpm
```

4. Switch to MDM single mode, by running the following commands on MDM 1:

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

```
scli --start_upgrade
```

```
scli --allow_commands_during_upgrade
```

```
scli --switch_cluster_mode --cluster_mode 1_node --
remove_slave_mdm_ip <MDM2 IP> --remove_tb_ip <TB1 IP>
```

In the `--switch_cluster_mode` command, type the data (not management) IP addresses.

5. Upgrade MDM 2 and TB 1, by running the following commands on both:
   a. Stop service:

```
/opt/emc/scaleio/mdm/bin/delete_service.sh
```

   b. Upgrade RPM:

```
rpm -U EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm
```

   c. Start service:

```
/opt/emc/scaleio/mdm/bin/create_service.sh
```

6. Add the upgraded MDM 2 and TB 1, by running the following commands on MDM 1:
   a. Add the components to the cluster:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_ip <MDM2 IP> --add_tb_ip <TB1 IP>
```

   b. Verify the cluster is in normal state:

```
scli --query_cluster
```

7. Switch MDM ownership, by running the following command on MDM 1:

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM2 IP>
```

8. Run the following commands on MDM 2:

a. Log in:

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

b. Enable commands:

```
scli --allow_commands_during_upgrade
```

c. View cluster information:

```
scli --query_cluster
```

d. Switch to single mode:

```
scli --switch_cluster_mode --cluster_mode 1_node --
remove_slave_mdm_ip <MDM1 IP> --remove_tb_ip <TB1 IP>
```

9. Upgrade MDM 1, by running the following commands on MDM 1:

a. Stop service:

```
/opt/emc/scaleio/mdm/bin/delete_service.sh
```

b. Upgrade RPM:

```
rpm -U EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm
```

c. Start service:

```
/opt/emc/scaleio/mdm/bin/create_service.sh
```

10. Add the MDM as a standby MDM, and switch to 3-node cluster mode, by running these commands from the Master MDM (MDM 2):

a. Log in (if not logged in):

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

b. Enable commands (if not already done):

```
scli --allow_commands_during_upgrade
```

c. Switch mode:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_ip <MDM1 IP> --add_tb_ip <TB1 IP>
```

d. Wait for the cluster to finish synching the repository:

```
scli --query_cluster
```

Wait until cluster state is Normal.

e. Switch cluster ownership:

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM1 IP>
```

11. Display list of all SDS, by running the following commands from Master MDM 1:

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

```
scli --query_all_sds
```

You can upgrade all SDSs within the same Fault Set.

Note

Between the upgrade of each set of SDSs, wait until all rebuilds are complete, and five seconds longer.

12. Upgrade all SDSs, one-at-a-time:

a. Enter an SDS into maintenance mode, by running the following command from MDM 1:

```
scli --enter_maintenance_mode --sds_id <SDS ID>
```

Validate that the operation succeeded.

b. Disable RFcache (if installed), by running the following command from MDM 1:

```
scli --disable_sds_rfcache --sds_id <SDS ID>
```

c. Upgrade the SDS, by running the following command from the SDS:

```
rpm -U EMC-ScaleIO-sds-2.0-13000.X.<flavor>.x86_64.rpm
```

d. Upgrade RFcache (if installed), by running the following command from the SDS:

```
rpm -U EMC-ScaleIO-xcache-2.0-13000.X.<flavor>.x86_64.rpm
```

e. Enable RFcache (if installed), by running the following command from MDM 1:

```
scli --enable_sds_rfcache --sds_id <SDS ID>
```

f. Remove the SDS from maintenance mode, by running the following command from MDM 1:

```
scli --exit_maintenance_mode --sds_id <SDS ID>
```

g. Before continuing, validate that the SDS is in the proper mode, by running the following command from MDM 1:

```
scli --query_properties --object_type SDS --object_id <SDS
ID> --properties MAINTENANCE_MODE_STATE
```

Ensure that the state is NO_MAINTENANCE

The SDS is upgraded. Repeat these steps for every SDS.

13. Upgrade all SDCs:

```
rpm -U EMC-ScaleIO-sdc-2.0-13000.X.<flavor>.x86_64.rpm
```

If the SDC kernel module is processing IO, its upgrade will occur only after the server is restarted.

The SDC is upgraded. Repeat this step for every SDC.

14. Validate the upgrade, by running the following command on MDM 1:

```
scli --query_upgrade
```

Ensure that the output includes the following:

```
0 SDSs need to be upgraded in total
```

15. Finalize the upgrade, by running the following command on MDM 1:

```
scli --finalize_upgrade
```

Ensure that the output includes the following:

```
Upgrade State: No Upgrade
```

16. Upgrade the currently installed Gateway by running the following command (optional):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_40/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

NOTICE

If your gateway uses a non-default certificate, you must copy the certificate and other files before upgrading the gateway. For more information, see "Upgrading the Gateway when a non-default certificate is used".

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 148.

After upgrading

After the upgrade, you can perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses

These options, and others, are described in "Maintaining a ScaleIO system".

Manual upgrade of Linux systems from v2.0.0.2+, in a 5-node cluster

Manual upgrade of ScaleIO v2.0.0.2 (or later) components, on a Linux server (all flavors) when the MDM cluster is in 5-node mode.

Before you begin

- Ensure that there is at least 1.1 GB of available disk space on all MDMs (additional 0.1 GB is required if the server also has an SDS).
- Java 1.8 is required on the Gateway and the GUI servers before performing the upgrade. Download the latest JRE for your operating system from the Oracle website.
- If needed, install the v2.0.1 `GPG-RPM-KEY` file (with the `rpm --import <path_to_new_public_key>` command). This is used to used to authenticate the RPM packages.

- To upgrade Ubuntu or CoreOS servers, you must first extract all the packages, as described in Extracting ScaleIO packages on page 160.
- Run the ScaleIO system analysis to analyze the ScaleIO system immediately prior to upgrading. It will identify and avoid faulty setups in the system. For more information, see System analysis overview on page 148.

When SNMP is enabled in the 2.x base version, SNMP MDM credentials remain in the v2.x Lockbox after the upgrade.

When performing this procedure on Xen server, use `siocli` in place of `scli`.

For the purposes of this procedure, we will use the following nomenclature:

MDM1—Master MDM

MDM2—Slave1 MDM

MDM3—Slave2 MDM

TB1—Tie Breaker1

TB2—Tie Breaker2

Procedure

1. Download and extract the ScaleIO packages to `/tmp/`.
2. Upgrade the GUI (on the Windows server):

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

3. Upgrade (or install) LIA on all machines in the system:

```
rpm -U EMC-ScaleIO-lia-2.0-13000.X.<flavor>.x86_64.rpm
```

4. Start the upgrade, by running the following commands on MDM 1:

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

```
scli --start_upgrade
```

5. Upgrade the MDM on the servers (except for the Master MDM).

   Perform the following steps on each server, one-at-a-time, in the following order:

   a. MDM2
   b. MDM3
   c. TB1
   d. TB2

   After validating MDM2, repeat this step for MDM3, and then the following servers.

   a. Stop service:

```
/opt/emc/scaleio/mdm/bin/delete_service.sh
```

b. Upgrade RPM:

```
rpm -U EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm
```

c. Start service:

```
/opt/emc/scaleio/mdm/bin/create_service.sh
```

d. Validate the cluster is in normal state:

```
scli --query_cluster
```

e. Repeat this group of steps, for each server.

The MDM packages are upgraded, per component.

6. Switch MDM ownership, by running the following commands on MDM1:

```
scli --allow_commands_during_upgrade
```

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM2 IP>
```

7. Upgrade MDM1, by running the following commands on it:

a. Stop service:

```
/opt/emc/scaleio/mdm/bin/delete_service.sh
```

b. Upgrade RPM:

```
rpm -U EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm
```

c. Start service:

```
/opt/emc/scaleio/mdm/bin/create_service.sh
```

8. Wait for the cluster to complete the repository synchronization, by running the following command on MDM1:

```
scli --query_cluster
```

Wait until the cluster state is normal.

9. Switch MDM ownership, by running the following commands on MDM 2:

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

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM1 IP>
```

10. Display list of all SDS, by running the following commands from Master MDM 1:

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

```
scli --query_all_sds
```

You can upgrade all SDSs within the same Fault Set.

Note

Between the upgrade of each set of SDSs, wait until all rebuilds are complete, and five seconds longer.

11. Upgrade all SDSs, one-at-a-time:
    a. Enter an SDS into maintenance mode, by running the following command from MDM 1:

    ```
    scli --enter_maintenance_mode --sds_id <SDS ID>
    ```

    Validate that the operation succeeded.

    b. Disable RFcache (if installed), by running the following command from MDM 1:

    ```
    scli --disable_sds_rfcache --sds_id <SDS ID>
    ```

    c. Upgrade the SDS, by running the following command from the SDS:

    ```
    rpm -U EMC-ScaleIO-sds-2.0-13000.X.<flavor>.x86_64.rpm
    ```

    d. Upgrade RFcache (if installed), by running the following command from the SDS:

    ```
    rpm -U EMC-ScaleIO-xcache-2.0-13000.X.<flavor>.x86_64.rpm
    ```

    e. Enable RFcache (if installed), by running the following command from MDM 1:

    ```
    scli --enable_sds_rfcache --sds_id <SDS ID>
    ```

f. Remove the SDS from maintenance mode, by running the following command from MDM 1:

```
scli --exit_maintenance_mode --sds_id <SDS ID>
```

g. Before continuing, validate that the SDS is in the proper mode, by running the following command from MDM 1:

```
scli --query_properties --object_type SDS --object_id <SDS
ID> --properties MAINTENANCE_MODE_STATE
```

Ensure that the state is NO_MAINTENANCE

The SDS is upgraded. Repeat these steps for every SDS.

12. Upgrade all SDCs:

```
rpm -U EMC-ScaleIO-sdc-2.0-13000.X.<flavor>.x86_64.rpm
```

If the SDC kernel module is processing IO, its upgrade will occur only after the server is restarted.

The SDC is upgraded. Repeat this step for every SDC.

13. Validate the upgrade, by running the following command on MDM 1:

```
scli --query_upgrade
```

Ensure that the output includes the following:

```
0 SDSs need to be upgraded in total
```

14. Finalize the upgrade, by running the following command on MDM 1:

```
scli --finalize_upgrade
```

Ensure that the output includes the following:

```
Upgrade State: No Upgrade
```

15. Upgrade the currently installed Gateway by running the following command (optional):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_40/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

NOTICE

If your gateway uses a non-default certificate, you must copy the certificate and other files before upgrading the gateway. For more information, see "Upgrading the Gateway when a non-default certificate is used".

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 148.

After upgrading

After the upgrade, you can perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses

These options, and others, are described in "Maintaining a ScaleIO system".

CHAPTER 7

Manual Upgrade of Windows Server

This chapter describes how to perform a manual upgrade of ScaleIO components on a Windows server. Topics include:

- Manual upgrade of physical Windows configurations 96
- Manual upgrade of Windows server from v1.32.x 96
- Manual upgrade of Windows server from v2.0 or v2.0.0.1, in a 3-node cluster .. 101
- Manual upgrade of Windows server from v2.0 or v2.0.0.1, in a 5-node cluster 101
- Manual upgrade of Windows server from v2.0.0.2+, in a 3-node cluster 102
- Manual upgrade of Windows server from v2.0.0.2+, in a 5-node cluster 106

Manual upgrade of physical Windows configurations

This section describes how to perform a manual upgrade when all ScaleIO components reside on physical Windows servers.

It is highly recommended to perform upgrades using the Installation Manager.

Note

When performing a manual upgrade, it is crucial to follow the described upgrade procedures step by step. Failure to do so might result in data loss and system instability.

Upgrading is non-disruptive; you can upgrade while IOs are running and volumes are mapped.

Extending from a 3-node to a 5-node cluster is performed, and described, after upgrading.

Continue to the section that matches your system environment.

Manual upgrade of Windows server from v1.32.x

Manual upgrade of ScaleIO v1.32.4 (or later) components, on a Windows server. To upgrade from pre-1.32.4 versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.

Before you begin

Ensure that there is at least 1.1 GB of available disk space on all MDMs (additional 0.1 GB is required if the server also has an SDS).

Java 1.8 is required on the Gateway and the GUI servers before performing the upgrade. Download the latest JRE for your operating system from the Oracle website.

When SNMP is enabled in the 1.32.x base version, the previous SNMP MDM credentials are erased from the old `gatewayUser.properties` configuration file and moved into the v2.x Lockbox. To harden the password, you can use the SioGWTool after the upgrade, or you can pass an environment variable while upgrading the Gateway. For more information, see "Non-Disruptive Upgrade Issues" in the *ScaleIO Upgrade Guide*.

When upgrading from v1.32.x, do not delete volumes or snapshots during the MDM cluster upgrade phase of the upgrade. This phase is normally very short.

For the purposes of this procedure, we will use the following nomenclature:

MDM 1—Primary MDM

MDM 2—Secondary MDM

TB 1—TieBreaker

Procedure

1. Download and extract the ScaleIO packages.
2. Upgrade the GUI, by running this file on the server where the GUI is running:

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

3. Use the Windows **Add/Remove Programs** to remove the v1.32 CallHome component from MDM servers.
4. Upgrade (or install) LIA on all machines in the system, by running this file:

```
EMC-ScaleIO-lia-2.0-13000.X.msi
```

5. Switch to MDM single mode, by running the following commands on MDM 1:

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

```
scli --switch_to_single_mode
```

6. Remove the Secondary MDM and the TieBreaker, by running the following commands on MDM 1:

```
scli --remove_secondary_mdm
```

```
scli --remove_tb
```

7. Upgrade the Secondary MDM, by running the following commands on MDM 2:
   a. Stop the EMC meta-data manager service
   b. Set role, by adding the following line into the `C:\Program Files\EMC\scaleio\mdm\cfg\conf.txt file`:

      `actor_role_is_manager=1`

   c. Upgrade the MDM, by running this file:

```
EMC-ScaleIO-mdm-2.0-13000.X.msi
```

8. Upgrade the Tie-Breaker:
   a. Use the Control Panel to remove EMC-scaleio-tb.
   b. Install the new Tie-Breaker, by running this file:

```
EMC-ScaleIO-mdm-2.0-13000.X.msi
```

9. Add the upgraded MDM 2 and TB1, by running the following commands on MDM 1:
   a. Add the Secondary MDM:

```
scli --add_secondary_mdm --secondary_mdm_ip <MDM2 IP>
```

b. Add the TB:

```
scli --add_tb --tb_ip <TB1 IP>
```

10. Switch to cluster mode, by running the following commands on MDM 1:

```
scli --switch_to_cluster_mode
```

Wait for the repository sync to complete, indicated by its state being normal.

```
scli --query_cluster
```

11. Switch MDM ownership, by running the following command on MDM 1:

```
scli --switch_mdm_ownership
```

12. Modify the new Master MDM (MDM2), by running the following commands on MDM 2:

a. Log in:

```
scli --login --username <USER> --password <PASSWORD> --
use_nonsecure_communication
```

b. Enable commands:

```
scli --allow_commands_during_upgrade --
use_nonsecure_communication
```

c. View cluster information:

```
scli --query_cluster --use_nonsecure_communication
```

d. Switch to single mode:

```
scli --switch_cluster_mode --cluster_mode 1_node --
remove_slave_mdm_ip <MDM1 IP> --remove_tb_ip <TB1 IP> --
use_nonsecure_communication
```

e. Remove the standby MDM:

```
scli --remove_standby_mdm --remove_mdm_ip <MDM1 IP> --
use_nonsecure_communication
```

f. Modify the management IP address:

```
scli --modify_management_ip --new_mdm_management_ip <MDM2
management IP> --target_mdm_ip <MDM2 IP> --
allow_duplicate_management_ips --use_nonsecure_communication
```

13. Upgrade MDM1, by running the following commands on it:

    a. Stop the EMC meta-data manager service

    b. Set role, by adding the following line into the `C:\Program Files\EMC\scaleio\mdm\cfg\conf.txt file`:

    `actor_role_is_manager=1`

    c. Upgrade the MDM, by running this file:

```
EMC-ScaleIO-mdm-2.0-13000.X.msi
```

14. Add the MDM as a standby MDM, by running these commands from the Master MDM (MDM 2):

    a. Log in (if not logged in):

```
scli --login --username <USER> --password <PASSWORD> --
use_nonsecure_communication
```

    b. Enable commands (if not already done):

```
scli --allow_commands_during_upgrade --
use_nonsecure_communication
```

    c. Add the standby MDM:

```
scli --add_standby_mdm --new_mdm_ip <MDM1 IP> --mdm_role
manager --new_mdm_management_ip <MDM1 management IP> --
new_mdm_name <MDM1 name> --use_nonsecure_communication
```

15. Switch to 3-node cluster mode, by running the following commands, from MDM 2:

    a. Switch mode:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_ip <MDM1 IP> --add_tb_ip <TB1 IP> --
use_nonsecure_communication
```

    b. View the cluster status:

```
scli --query_cluster --use_nonsecure_communication
```

Wait until cluster state is Normal, which indicates that for the repository synching is complete.

c. Switch cluster ownership:

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM1 IP> --
use_nonsecure_communication
```

16. Upgrade all SDSs, by running the following file on each SDS server:

```
EMC-ScaleIO-sds-2.0-13000.X.msi
```

You can upgrade (and reboot) all SDSs in the same Fault Set or Protection Domain. You can upgrade (and reboot) all SDSs or Fault Sets in the same Protection Domain.

Note

Between the upgrade of each set of SDSs, wait until all rebuilds are complete, and five seconds longer.

17. Upgrade all SDCs, by running the following file on each SDC server:

```
EMC-ScaleIO-sdc-2.0-13000.X.msi
```

18. Restart the SDC server.
19. Validate the upgrade, by running the following command on MDM 1:

```
scli --query_upgrade --use_nonsecure_communication
```

20. Finalize the upgrade, by running the following command on MDM 1:

```
scli --finalize_upgrade --use_nonsecure_communication
```

21. Upgrade the currently installed Gateway by running the following file (optional):

```
EMC-ScaleIO-gateway-2.0-13000.X-x64.msi
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

NOTICE

If your gateway uses a non-default certificate, before upgrading the gateway you must copy the certificate and other files. For more information, see "Upgrading the Gateway when a non-default certificate is used".

After upgrading

These upgrade procedures leave the system in non-secure authentication mode. After the upgrade, you can switch to secure mode, as well as perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node

These options, and others, are described in "Maintaining a ScaleIO system".

Note

If you choose not to switch to secure mode, you need to disable secure communication on every MDM server before you can run SCLI commands.

Manual upgrade of Windows server from v2.0 or v2.0.0.1, in a 3-node cluster

Manual upgrade of ScaleIO v2.0 or v2.0.0.1 components, on a Windows server. This procedure is valid when the MDM cluster is in 3-node mode.

This procedure is completed as follows:

- First, upgrade from the base version to v2.0.0.x (latest), using the same steps described in Manual upgrade of Windows server from v2.0.0.2+, in a 3-node cluster on page 102.
  In the upload packages step, upload the base version and the v2.0.0.x (latest) packages. These packages are of the following format: `EMC-ScaleIO-<component>-2.0-XXXX.X.msi`
- Then, upgrade from v2.0.0.x (latest) to the current version, using those exact same steps.
  In the upload packages step, upload the v2.0.0.x (latest) and the current version packages. These packages are of the following format:
  - v2.0.0.x (latest) packages: `EMC-ScaleIO-<component>-2.0-XXXX.X.msi`
  - Current version packages: `EMC-ScaleIO-<component>-2.0-13000.X.msi`

Manual upgrade of Windows server from v2.0 or v2.0.0.1, in a 5-node cluster

Manual upgrade of ScaleIO v2.0 or v2.0.0.1 components, on a Windows server. This procedure is valid when the MDM cluster is in 5-node mode.

This procedure is completed as follows:

- First, upgrade from the base version to v2.0.0.x (latest), using the same steps described in Manual upgrade of Windows server from v2.0.0.2+, in a 5-node cluster on page 106.
  In the upload packages step, upload the base version and the v2.0.0.x (latest) packages. These packages are of the following format: `EMC-ScaleIO-<component>-2.0-XXXX.X.msi`

In the upload packages step, upload the base version and the v2.0.0.x (latest) packages.

- Then, upgrade from v2.0.0.x (latest) to the current version, using those exact same steps.
  In the upload packages step, upload the v2.0.0.x (latest) and the current version packages. These packages are of the following format:
  - v2.0.0.x (latest) packages: `EMC-ScaleIO-<component>-2.0-XXXX.X.msi`
  - Current version packages: `EMC-ScaleIO-<component>-2.0-13000.X.msi`

Manual upgrade of Windows server from v2.0.0.2+, in a 3-node cluster

Manual upgrade of ScaleIO v2.0.0.2 (or later) components, on a Windows server. This procedure is valid when the MDM cluster is in 3-node mode.

Before you begin

Ensure that there is at least 1.1 GB of available disk space on all MDMs (additional 0.1 GB is required if the server also has an SDS).

Java 1.8 is required on the Gateway and the GUI servers before performing the upgrade. Download the latest JRE for your operating system from the Oracle website.

If the Windows OS is installed on the SATADOM, you must make changes to the memory settings before upgrading, as described in Upgrading Windows servers when the OS is installed on SATADOM on page 176.

When SNMP is enabled in the 2.x base version, SNMP MDM credentials remain in the v2.x Lockbox after the upgrade.

For the purposes of this procedure, we will use the following nomenclature:

MDM 1—Primary MDM

MDM 2—Secondary MDM

TB 1—Tie Breaker

Procedure

1. Download and extract the ScaleIO packages to a temporary folder.
2. Upgrade the GUI (on the Windows server):

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

3. Upgrade (or install) LIA on all machines in the system:

```
EMC-ScaleIO-lia-2.0-13000.X.msi
```

4. Switch to MDM single mode, by running the following commands on MDM 1:

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

```
scli --start_upgrade
```

```
scli --allow_commands_during_upgrade
```

```
scli --switch_cluster_mode --cluster_mode 1_node --
remove_slave_mdm_ip <MDM2 IP> --remove_tb_ip <TB1 IP>
```

In the `--switch_cluster_mode` command, type the data (not management) IP addresses.

5. Upgrade MDM 2 and TB 1, by running this file on both servers:

```
EMC-ScaleIO-mdm-2.0-13000.X.msi
```

6. Add the upgraded MDM 2 and TB 1, by running the following commands on MDM 1:

   a. Add the components to the cluster:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_ip <MDM2 IP> --add_tb_ip <TB1 IP>
```

   b. Verify the cluster is in normal state:

```
scli --query_cluster
```

7. Switch MDM ownership, by running the following command on MDM 1:

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM2 IP>
```

8. Run the following commands on MDM 2:

   a. Log in:

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

   b. Enable commands:

```
scli --allow_commands_during_upgrade
```

c. View cluster information:

```
scli --query_cluster
```

d. Switch to single mode:

```
scli --switch_cluster_mode --cluster_mode 1_node --
remove_slave_mdm_ip <MDM1 IP> --remove_tb_ip <TB1 IP>
```

9. Upgrade MDM 1, by running this file on it:

```
EMC-ScaleIO-mdm-2.0-13000.X.msi
```

10. Add the MDM as a standby MDM, and switch to 3-node cluster mode, by running these commands from the Master MDM (MDM 2):

a. Log in (if not logged in):

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

b. Enable commands (if not already done):

```
scli --allow_commands_during_upgrade
```

c. Switch mode:

```
scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_ip <MDM1 IP> --add_tb_ip <TB1 IP>
```

d. Wait for the cluster to finish synching the repository:

```
scli --query_cluster
```

Wait until cluster state is Normal.

e. Switch cluster ownership:

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM1 IP>
```

11. Display list of all SDS, by running the following commands from Master MDM 1:

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

```
scli --query_all_sds
```

You can upgrade all SDSs within the same Fault Set.

Note

Between the upgrade of each set of SDSs, wait until all rebuilds are complete, and five seconds longer.

12. Upgrade all SDSs:
    a. Enter an SDS into maintenance mode, by running the following command from MDM 1:

    ```
    scli --enter_maintenance_mode --sds_id <SDS ID>
    ```

    Validate that the operation succeeded.

    b. Disable RFcache (if installed), by running the following command from MDM 1:

    ```
    scli --disable_sds_rfcache --sds_id <SDS ID>
    ```

    c. Upgrade the SDS, by running from the SDS:

    ```
    EMC-ScaleIO-sds-2.0-13000.X.msi
    ```

    d. Upgrade and enable RFcache (if installed):
        a. From the SDS, run this file:

        ```
        EMC-ScaleIO-xcache-2.0-13000.X.msi
        ```

        b. Enable RFcache, by running the following command from MDM 1:

        ```
        scli --enable_sds_rfcache --sds_id <SDS ID>
        ```

        c. Restart the SDS server.

    e. Remove the SDS from maintenance mode, by running the following command from MDM 1:

    ```
    scli --exit_maintenance_mode --sds_id <SDS ID>
    ```

    f. Before continuing, validate that the SDS is in the proper mode, by running the following command from MDM 1:

    ```
    scli --query_properties --object_type SDS --object_id <SDS
    ID> --properties MAINTENANCE_MODE_STATE
    ```

    Ensure that the state is NO_MAINTENANCE

The SDS is upgraded. Repeat this upgrade step for every SDS.

13. Upgrade all SDCs:
    a. Upgrade the SDC:

    ```
    EMC-ScaleIO-sdc-2.0-13000.X.msi
    ```

    b. Restart the SDC server.

    The SDC is upgraded. Repeat these steps for every SDC.
14. Validate the upgrade, by running the following command on MDM 1:

    ```
    scli --query_upgrade
    ```

    Ensure that the output includes the following:

    ```
    0 SDSs need to be upgraded in total
    ```
15. Finalize the upgrade, by running the following command on MDM 1:

    ```
    scli --finalize_upgrade
    ```

    Ensure that the output includes the following:

    ```
    Upgrade State: No Upgrade
    ```
16. Upgrade the currently installed Gateway by running the following file (optional):

    ```
    EMC-ScaleIO-gateway-2.0-13000.X-x64.msi
    ```

    A success message is displayed, and the IM interface now reflects the updated Gateway version.

    NOTICE

    If your gateway uses a non-default certificate, before upgrading the gateway you must copy the certificate and other files. For more information, see "Upgrading the Gateway when a non-default certificate is used".

Manual upgrade of Windows server from v2.0.0.2+, in a 5-node cluster

Manual upgrade of ScaleIO v2.0.0.2 (or later) components, on a Windows server. This procedure is valid when the MDM cluster is in 5-node mode.

Before you begin

Ensure that there is at least 1.1 GB of available disk space on all MDMs (additional 0.1 GB is required if the server also has an SDS).

Java 1.8 is required on the Gateway and the GUI servers before performing the upgrade. Download the latest JRE for your operating system from the Oracle website.

If the Windows OS is installed on the SATADOM, you must make changes to the memory settings before upgrading, as described in Upgrading Windows servers when the OS is installed on SATADOM on page 176.

When SNMP is enabled in the 2.x base version, SNMP MDM credentials remain in the v2.x Lockbox after the upgrade.

For the purposes of this procedure, we will use the following nomenclature:

MDM1—Master MDM

MDM2—Slave1 MDM

MDM3—Slave2 MDM

TB1—Tie Breaker1

TB2—Tie Breaker2

Procedure

1. Download and extract the ScaleIO packages to a temporary folder.
2. Upgrade the GUI (on the Windows server):

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

3. Upgrade (or install) LIA on all machines in the system:

```
EMC-ScaleIO-lia-2.0-13000.X.msi
```

4. Start the upgrade, by running the following commands on MDM 1:

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

```
scli --start_upgrade
```

5. Upgrade the MDM on the servers (except for the Master MDM).

   Perform the following steps on each server, one-at-a-time, in the following order:

   a. MDM2
   b. MDM3
   c. TB1
   d. TB2

   After validating MDM2, repeat this step for MDM3, and then the following servers.

   a. Upgrade the MDM component:

```
EMC-ScaleIO-mdm-2.0-13000.X.msi
```

b. Validate the cluster is in normal state:

```
scli --query_cluster
```

c. Repeat this group of steps, for each server.

The MDM packages are upgraded.

6. Switch MDM ownership, by running the following commands on MDM1:

```
scli --allow_commands_during_upgrade
```

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM2 IP>
```

7. Upgrade MDM1, by running the following file on it:

```
EMC-ScaleIO-mdm-2.0-13000.X.msi
```

8. Wait for the cluster to complete the repository synchronization, by running the following command on MDM2:

```
scli --query_cluster
```

Wait until the cluster state is normal.

9. Switch MDM ownership, by running the following commands on MDM 2:

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

```
scli --switch_mdm_ownership --new_master_mdm_ip <MDM1 IP>
```

10. Display list of all SDS, by running the following commands from Master MDM 1:

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

```
scli --query_all_sds
```

You can upgrade all SDSs within the same Fault Set.

Note

Between the upgrade of each set of SDSs, wait until all rebuilds are complete, and five seconds longer.

11. Upgrade all SDSs:

a. Enter an SDS into maintenance mode, by running the following command from MDM 1:

```
scli --enter_maintenance_mode --sds_id <SDS ID>
```

Validate that the operation succeeded.

b. Disable RFcache (if installed), by running the following command from MDM 1:

```
scli --disable_sds_rfcache --sds_id <SDS ID>
```

c. Upgrade the SDS, by running from the SDS:

```
EMC-ScaleIO-sds-2.0-13000.X.msi
```

d. Upgrade and enable RFcache (if installed):

a. From the SDS, run this file:

```
EMC-ScaleIO-xcache-2.0-13000.X.msi
```

b. Enable RFcache, by running the following command from MDM 1:

```
scli --enable_sds_rfcache --sds_id <SDS ID>
```

c. Restart the SDS server.

e. Remove the SDS from maintenance mode, by running the following command from MDM 1:

```
scli --exit_maintenance_mode --sds_id <SDS ID>
```

f. Before continuing, validate that the SDS is in the proper mode, by running the following command from MDM 1:

```
scli --query_properties --object_type SDS --object_id <SDS
ID> --properties MAINTENANCE_MODE_STATE
```

Ensure that the state is NO_MAINTENANCE

The SDS is upgraded. Repeat this upgrade step for every SDS.

12. Upgrade all SDCs:

a. Upgrade the SDC:

```
EMC-ScaleIO-sdc-2.0-13000.X.msi
```

b. Restart the SDC server.

The SDC is upgraded. Repeat these steps for every SDC.

13. Validate the upgrade, by running the following command on MDM 1:

```
scli --query_upgrade
```

Ensure that the output includes the following:

```
0 SDSs need to be upgraded in total
```

14. Finalize the upgrade, by running the following command on MDM 1:

```
scli --finalize_upgrade
```

Ensure that the output includes the following:

```
Upgrade State: No Upgrade
```

15. Upgrade the currently installed Gateway by running the following file (optional):

```
EMC-ScaleIO-gateway-2.0-13000.X-x64.msi
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

NOTICE

If your gateway uses a non-default certificate, before upgrading the gateway you must copy the certificate and other files. For more information, see "Upgrading the Gateway when a non-default certificate is used".

CHAPTER 8

Manual Upgrade of Xen Server

This chapter describes how to perform a manual upgrade of ScaleIO components on a Xen server. Topics include:

- Manual upgrade of Xen system from v1.32.x 112

Manual upgrade of Xen system from v1.32.x

Manual upgrade of ScaleIO v1.32.4 (or later) components, on a Xen server. To upgrade from pre-1.32.4 versions, you must first upgrade to the latest 1.32.x version, then proceed with the upgrade described here.

Before you begin

Ensure that there is at least 1.1 GB of available disk space on all MDMs (additional 0.1 GB is required if the server also has an SDS).

If needed, install the v2.0.1 `GPG-RPM-KEY` file (with the `rpm --import <path_to_new_public_key>` command). This is used to used to authenticate the RPM packages.

Java 1.8 is required on the Gateway and the GUI servers before performing the upgrade. Download the latest JRE for your operating system from the Oracle website.

When upgrading from v1.32.x, do not delete volumes or snapshots during the MDM cluster upgrade phase of the upgrade. This phase is normally very short.

For the purposes of this procedure, we will use the following nomenclature:

MDM 1—Primary MDM

MDM 2—Secondary MDM

MDM 3—TieBreaker

Procedure

1. Download and extract the ScaleIO packages to `/tmp/`.
2. Upgrade the GUI (on the Windows server):

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

3. Remove the v1.32 CallHome component, by running the following command on MDM 1 and MDM 2:

```
rpm -e `rpm -qa|grep ScaleIO|grep callhome`
```

4. Upgrade (or install) LIA on all machines in the system:

```
rpm -U EMC-ScaleIO-lia-2.0-13000.X.xs6.5.0.x86_64.rpm --
nosignature
```

5. Switch to MDM single mode, by running the following commands on MDM 1:

```
siocli --login --username <USER> --password <PASSWORD>
```

```
siocli --switch_to_single_mode
```

6. Remove the Secondary MDM and the TieBreaker, by running the following commands on MDM 1:

```
siocli --remove_secondary_mdm
```

```
siocli --remove_tb
```

7. Upgrade MDM 2, by running the following commands on it:
   a. Stop service:

   ```
   /opt/emc/scaleio/mdm/bin/delete_service.sh
   ```

   b. Set role:

   ```
   echo actor_role_is_manager=1 >> /opt/emc/scaleio/mdm/cfg/
   conf.txt
   ```

   c. Upgrade RPM:

   ```
   rpm -U EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm  --
   nosignature
   ```

   d. Start service:

   ```
   /opt/emc/scaleio/mdm/bin/create_service.sh
   ```

8. Upgrade the TieBreaker, by running the following commands on TB 1:
   a. Stop service:

   ```
   /opt/emc/scaleio/tb/bin/delete_service.sh
   ```

   b. Uninstall previous version:

   ```
   rpm -e --nosignature `rpm -qa --nosignature|grepScaleIO|
   grep tb`
   ```

   c. Install new version (with MDM, not TB package):

   ```
   rpm -i EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm --
   nosignature
   ```

9. Add the upgraded MDM 2 and TB 1, by running the following commands on MDM 1:

a. Add the Secondary MDM:

```
siocli --add_secondary_mdm --secondary_mdm_ip <MDM2 IP>
```

b. Add the TB:

```
siocli --add_tb --tb_ip <TB1 IP>
```

10. Switch to cluster mode, by running the following commands on MDM 1:

```
siocli --switch_to_cluster_mode
```

Wait for the repository sync to complete, indicated by its state being normal.

```
siocli --query_cluster
```

11. Switch MDM ownership, by running the following command on MDM 1:

```
siocli --switch_mdm_ownership
```

12. Modify the new Master MDM (MDM2), by running the following commands on MDM 2:

a. Log in:

```
siocli --login --username <USER> --password <PASSWORD> --
use_nonsecure_communication
```

b. Enable commands:

```
siocli --allow_commands_during_upgrade --
use_nonsecure_communication
```

c. View cluster information:

```
siocli --query_cluster --use_nonsecure_communication
```

d. Switch to single mode:

```
siocli --switch_cluster_mode --cluster_mode 1_node --
remove_slave_mdm_ip <MDM1 IP> --remove_tb_ip <TB1 IP> --
use_nonsecure_communication
```

e. Remove the standby MDM:

```
siocli --remove_standby_mdm --remove_mdm_ip <MDM1 IP> --
use_nonsecure_communication
```

f. Modify the management IP address (run the command on one line):

```
siocli --modify_management_ip --new_mdm_management_ip <MDM2
management IP> --target_mdm_ip <MDM2 IP>
--allow_duplicate_management_ips --
use_nonsecure_communication
```

13. Upgrade the MDM 1, by running the following commands on MDM 1:

a. Stop service:

```
/opt/emc/scaleio/mdm/bin/delete_service.sh
```

b. Set role:

```
echo actor_role_is_manager=1 >> /opt/emc/scaleio/mdm/cfg/
conf.txt
```

c. Upgrade RPM:

```
rpm -U EMC-ScaleIO-mdm-2.0-13000.X.<flavor>.x86_64.rpm --
nosignature
```

d. Start service:

```
/opt/emc/scaleio/mdm/bin/create_service.sh
```

14. Add the MDM as a standby MDM, by running these commands from the Master MDM (MDM 2):

a. Log in (if not logged in):

```
siocli --login --username <USER> --password <PASSWORD> --
use_nonsecure_communication
```

b. Enable commands (if not already done):

```
siocli --allow_commands_during_upgrade --
use_nonsecure_communication
```

c. Add the standby MDM:

```
siocli --add_standby_mdm --new_mdm_ip <MDM1 IP> --mdm_role
manager --new_mdm_management_ip <MDM1 management IP> --
new_mdm_name <MDM1 name> --use_nonsecure_communication
```

15. Switch to 3-node cluster mode, by running the following commands, from MDM 2:

a. Switch mode:

```
siocli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_ip <MDM1 IP> --add_tb_ip <TB1 IP> --
use_nonsecure_communication
```

b. Wait for the cluster to finish synching the repository:

```
siocli --query_cluster --use_nonsecure_communication
```

Wait until cluster state is Normal.

c. Switch cluster ownership:

```
siocli --switch_mdm_ownership --new_master_mdm_ip <MDM1 IP>
--use_nonsecure_communication
```

16. Upgrade all SDSs, by running the following command on each SDS server:

```
rpm -U EMC-ScaleIO-sds-2.0-13000.X.<flavor>.x86_64.rpm --
nosignature
```

You can update (and reboot) all SDSs in the same Fault Set. You can upgrade (and reboot) all SDSs or Fault Sets in the same Protection Domain.

Note

Between the upgrade of each set of SDSs, wait until all rebuilds are complete, and five seconds longer.

17. Upgrade all SDCs, by running the following commands on each SDC server:

```
rpm -U EMC-ScaleIO-sdc-2.0-13000.X.<flavor>.x86_64.rpm --
nosignature
```

If the SDC kernel module is processing IO, its upgrade will occur only after the server is restarted.

18. Validate the upgrade, by running the following command on MDM 1:

```
siocli --query_upgrade --use_nonsecure_communication
```

19. Finalize the upgrade, by running the following command on MDM 1:

```
siocli --finalize_upgrade --use_nonsecure_communication
```

20. Upgrade the currently installed Gateway by running the following command (optional):

```
SIO_GW_KEYTOOL=<keytool_path> rpm -U EMC-ScaleIO-
gateway-2.0-13000.X.x86_64.rpm (for SLES 11.3 add --nodeps)
```

Example:

```
SIO_GW_KEYTOOL=/usr/java/jre1.8.0_40/bin/ rpm -U /tmp/EMC-
ScaleIO-gateway-2.0-13000.X.x86_64.rpm (--nodeps)
```

A success message is displayed, and the IM interface now reflects the updated Gateway version.

NOTICE

If your gateway uses a non-default certificate, you must copy the certificate and other files before upgrading the gateway. For more information, see "Upgrading the Gateway when a non-default certificate is used".

After upgrading

After the upgrade, you can perform many other tasks, including:

- Install RFcache
- Extend the MDM cluster from 3-node to 5-node
- Add virtual IP addresses

These options, and others, are described in "Maintaining a ScaleIO system".

PART 4

Firmware Upgrade

This part describes how to upgrade firmware and drivers. Chapters include:

Chapter 9, "ScaleIO Ready Node Server Firmware Upgrades"

Chapter 10, "Upgrade of LSI RAID Controller Firmware and Driver"

CHAPTER 9

ScaleIO Ready Node Server Firmware Upgrades

This chapter explains how to upgrade the firmware and drivers on ScaleIO Ready Node servers. Topics include:

- Single server firmware and BIOS upgrade to a specific version.....122
- Multiple firmware upgrades on one or more servers.....127

Single server firmware and BIOS upgrade to a specific version

The procedures in this section provide instructions for upgrading the Dell firmware and BIOS to a specific version on a single ScaleIO Ready Node server.

Perform the procedures in the order in which they are presented.

Upgrading the Dell firmware and BIOS

When upgrading the firmware and BIOS, first upgrade the iDRAC firmware. After this is done, upgrade the BIOS and any other firmware that require a node reboot cycle.

Before you begin

You need a management server or service laptop from which to connect to the Dell Lifecyle Controller.

In the following example, the iDRAC firmware is upgraded from version 2.30.30.30 to 2.40.40.40, and the BIOS is upgraded from version 2.1.5 to 2.2.5.

Note

Use the BIOS upgrade and reboot instructions for any non-iDRAC firmware that requires an upgrade.

Procedure

1. From your Internet browser, go to https://*<iDRAC_IP_address>* and then log in.

   Use root as the username and password as the password.

   The iDRAC **Server Overview** page appears. The **Properties** tab displays the currently installed BIOS and firmware versions.

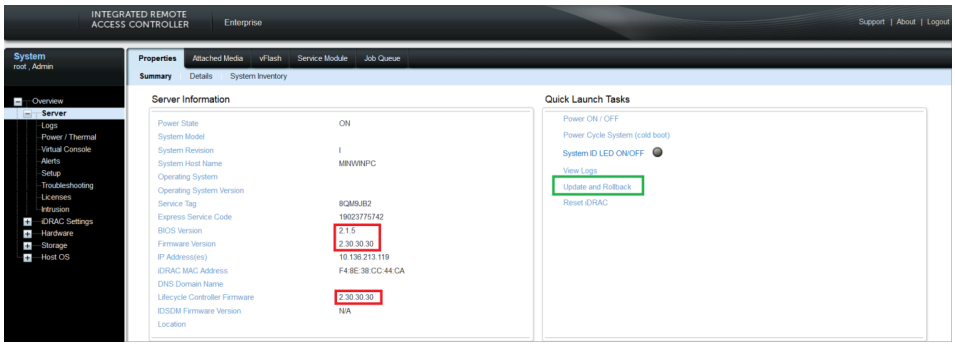

2. From the **Quick Launch** tasks, click **Update and Rollback**.

   The **Firmware** page appears:

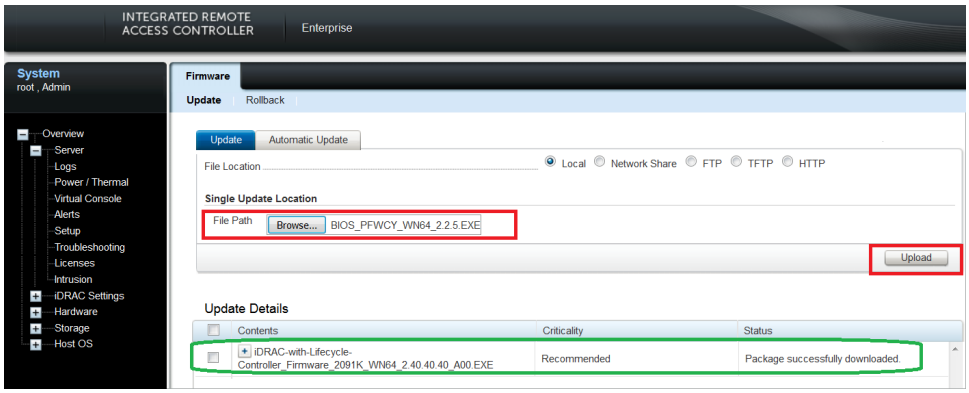

3. On the **Update** tab, click **Browse** and select the iDRAC firmware package—in this example, `iDRAC-with-Lifecycle-Controller_Firmware_2.40.40.40.exe`.
4. Click **Upload**.
5. Select the iDRAC firmware file and click **Install** to start the iDRAC firmware update:

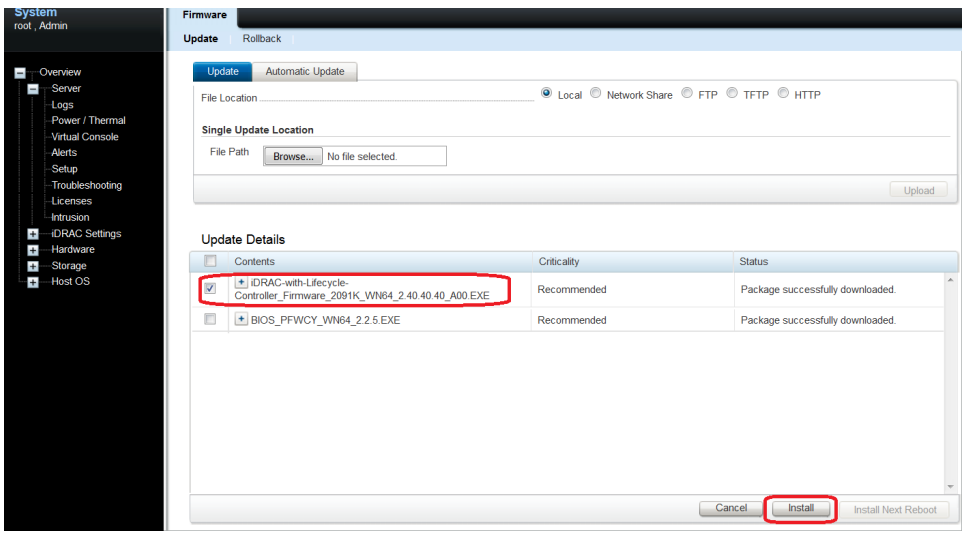

The iDRAC is updated and the web session is reset.

6. Reconnect to the iDRAC and log in.
7. In the **Server Overview** page, verify that the iDRAC controller was updated.

   The following example shows that the firmware was updated to version 2.40.40.40:

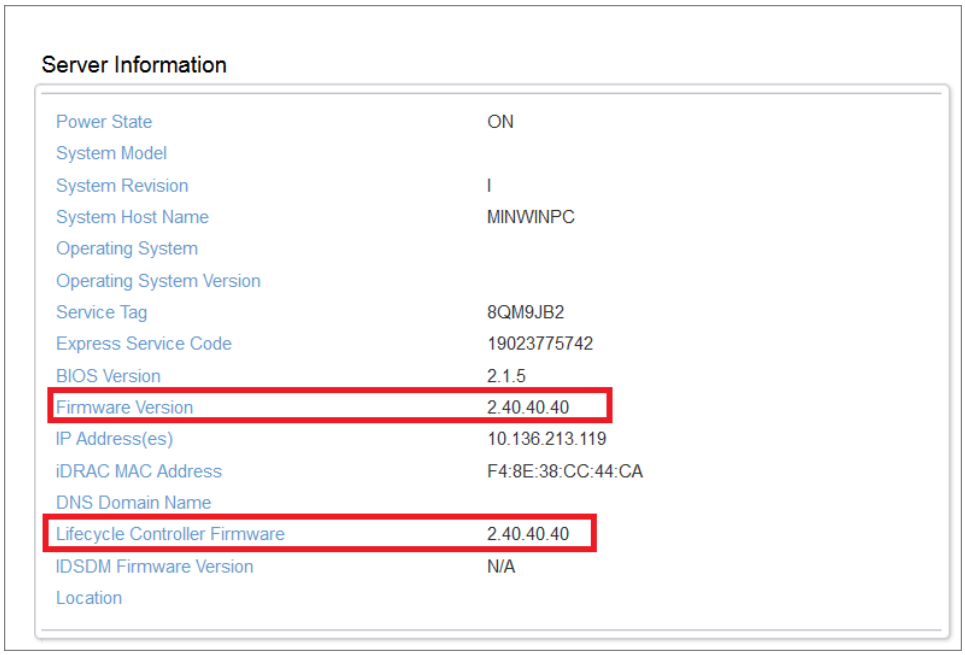

Note

If you are only updating the iDRAC firmware, there is no need to reboot the node. Instead, you can move to the next node to update the iDAC firmware, repeating steps 1 on page 122 - 7 on page 123 for each additional node.

8. To upgrade the BIOS or other firmware versions, launch a KVM session to the node.

Note

A reboot is required when upgrading the BIOS or other firmware, and the KVM session provides a window into the BIOS update process.

The Java applet downloads.

9. Enable the pop-up so that the SVM session can open.
10. On the iDRAC **Update** tab, click **Browse** and select the correct BIOS upgrade file—in this example, `BIOS_PFWCY_WN64_ 2.2.5.exe`.
11. Click **Upload**.
12. After the file has uploaded, click **Install Next Reboot**.

    The BIOS upgrade is queued for installation.
13. Repeat steps 10 on page 124 to 12 on page 124 for each additional firmware you want to upgrade.
14. If you are upgrading the BIOS or additional firmware, gracefully reboot the server, as described in Preparing the node for a graceful reboot on page 125.

Preparing the node for a graceful reboot

Prepare each node to reboot in a graceful fashion by first entering the node into maintenance mode.

Before you begin

Note

If you are only upgrading the iDRAC firmware, you are not required to reboot the server.

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

Procedure

1. Move all applications to a different node:
   - On an ESX node that is not a cluster member, and that is not configured for HA and DRS, migrate the VMs to another ESX.
   - 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.

     ---
2. Log in to the ScaleIO GUI as an admin user.
3. In ScaleIO **Backend** view, select **By SDSs** table view.

   Figure 1 Table view toggle options

   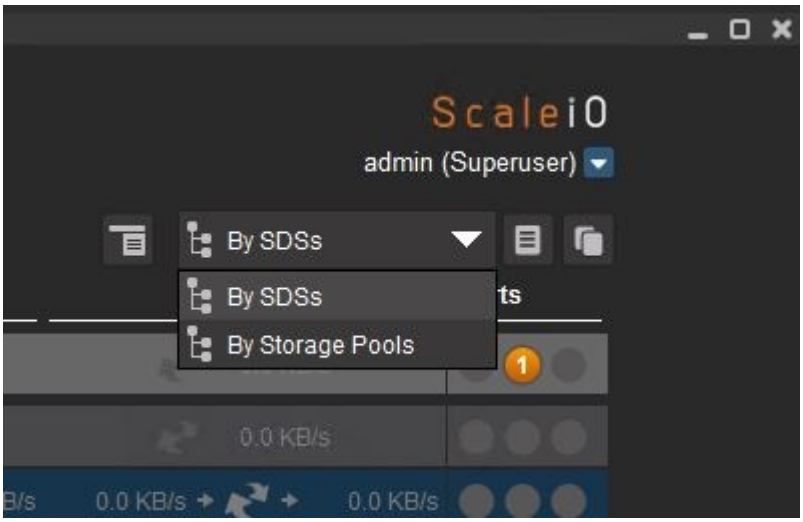

4. Right-click the SDS node you are rebooting, and select **Enter Maintenance Mode**.
5. In the **Enter maintenance mode** window, ensure that there are no errors, and then click **OK**.
6. When the operation finishes successfully, click **Close**.

   The node's IP address appears with a wrench next to it.

Figure 2 SDS displayed in maintenance mode

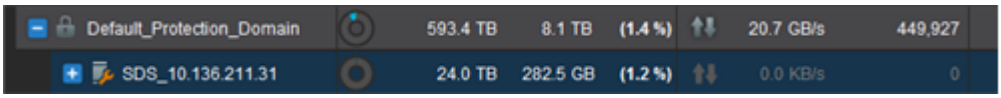

7. On an ESX node, enter maintenance mode:
   a. Log in to the vCenter via the vSphere Client or Web Client, and locate the relevant ESX IP address.
   b. Select the SVM, and from the **Basic Tasks** pane select **Shut down the virtual machine**.
   c. When the SVM is off, right-click the node and select **Enter Maintenance Mode**.
8. Obtain customer permission to reboot the node, and then gracefully 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 reboot after entering the SDS into maintenance mode.

Results

The firmware update will be performed automatically during the reboot cycle using the Dell Lifecycle Controller management utility.

Note

The node will reboot several times during installation of the new BIOS.

Return the node to operation

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

Procedure

1. Wait for the node to power on.

   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 Linux and Windows nodes, skip to the next step. On an ESX node:
   a. From the vSphere 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 **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 ESX nodes, you can migrate VMs to the node.

Results

The overview page on the iDRAC will indicate the latest version of the BIOS:

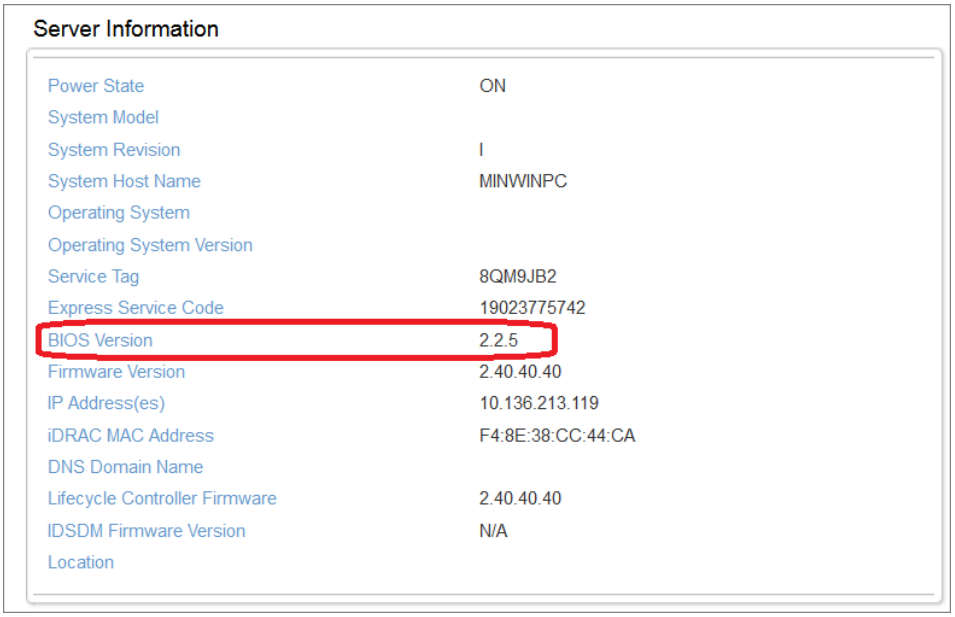

Server Information

Power State	ON
System Model	
System Revision	I
System Host Name	MINWINPC
Operating System	
Operating System Version	
Service Tag	8QM9JB2
Express Service Code	19023775742
BIOS Version	2.2.5
Firmware Version	2.40.40.40
IP Address(es)	10.136.213.119
iDRAC MAC Address	F4:8E:38:CC:44:CA
DNS Domain Name	
Lifecycle Controller Firmware	2.40.40.40
IDSDM Firmware Version	N/A
Location	

Update of the BIOS and the BMC is complete.

After you finish

Repeat the above procedures for updating the firmware and BIOS, gracefully rebooting the server, and returning it to operation on every node requiring a firmware upgrade.

Multiple firmware upgrades on one or more servers

The procedures in this section provide instructions for upgrading multiple Dell firmwares on one or more working ScaleIO Ready Node servers.

Perform the procedures in the order in which they are presented, for every server requiring a firmware update.

Opening the KVM console

Open the KVM console on a server in a ScaleIO Ready Node system.

Before you begin

Ensure that:

- The system environment meets the prerequisites for using the KVM console.
- You know the IP address of the BMC (iDRAC) port.
- You know the username and password for accessing the BMC (iDRAC) (default username/password are root/password).

Procedure

1. From your Internet browser, go to https://*<iDRAC_IP_address>*.
2. In the **DELL Console Login** window, type the user name and password, and click **Login**.
3. In the main KVM console window, select the **Server** node in the navigation pane.

   The **System Summary** dashboard is displayed.
4. In the **Virtual Console Preview** pane, click **Launch** to start a console session.

   A popup security warning screen is displayed.
5. Select **Accept** and click **Run**.

   The **Java Console** window opens and provides you with console access to the server.
   If this is the first time that you are opening a console, additional warning and confirmation prompts may appear. Click **OK** to grant approvals in these prompts.

   The subsequent steps depend on your browser selection and how it is configured. If downloads run automatically, the console window appears. If not, follow the instructions given in the next step to open the console window.
6. Depending on your browser, perform the following steps:

Browser	Steps
Internet Explorer	Click **Open** to run the `jviewer.jnlp` file.
Google Chrome	a. Click **Show all downloads**. b. In the **Downloads** window, at the warning message, click **Keep**. c. In the downloads list, click the downloaded file URL. d. At the warning message, click **Keep** or **Keep anyway**. e. Click the `jviewer.jnlp` file. If a security message appears, click **Run**. The management console window is displayed.

Results

The KVM console is open and ready for use.

Updating the BIOS, firmware and settings

ScaleIO Ready Node deployments require specific versions of drivers, BIOS, and firmware that have been qualified by Dell EMC. If the servers do not have the correct versions, you must update them.

A variety of factors can influence a mismatch between the required versions and the versions installed on the servers, such as firmware updates post server shipment, or a FRU replacement with a different firmware version than in the warehouse. For example, if you have replaced the server system board, the FRU's BIOs and iDRAC firmware versions will be different.

You are therefore required to verify that all server drivers, BIOS, and firmware meet the required versions, as described in Verifying status of system hardware, storage

controller, and disks on page 253, before deploying a server in the ScaleIO Ready Node environment.

To perform any updates needed to meet ScaleIO Ready Node requirements, use the ScaleIO Ready Node Hardware Update Bootable ISO ("Hardware ISO"). The Hardware ISO is based on the Dell OpenManage Deployment Toolkit (DTK). The DTK provides a framework of tools necessary for the configuration of PowerEdge servers. For ScaleIO, a custom script has been injected, along with specific qualified BIOS/ firmware update packages.

For additional information regarding the Hardware ISO, see the reference section, DTK - Hardware Update Bootable ISO on page 241.

Preparing the node for a graceful reboot

Prepare each node to reboot in a graceful fashion by first entering the node into maintenance mode.

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. Move all applications to a different node:
   - On an ESX node that is not a cluster member, and that is not configured for HA and DRS, migrate the VMs to another ESX.
   - 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.

2. Log in to the ScaleIO GUI as an admin user.
3. In ScaleIO **Backend** view, select **By SDSs** table view.

   Figure 3 Table view toggle options

   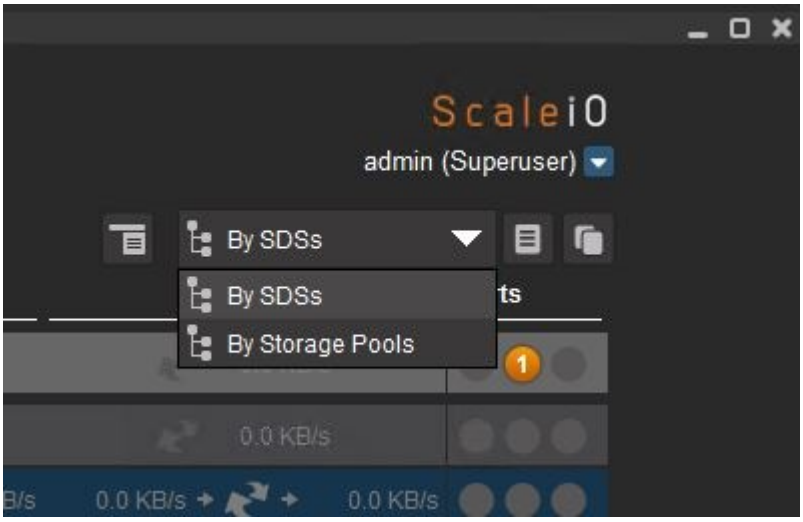

4. Right-click the SDS node you are rebooting, and select **Enter Maintenance Mode**.

5. In the **Enter maintenance mode** window, ensure that there are no errors, and then click **OK**.
6. When the operation finishes successfully, click **Close**.

   The node's IP address appears with a wrench next to it.

   Figure 4 SDS displayed in maintenance mode

   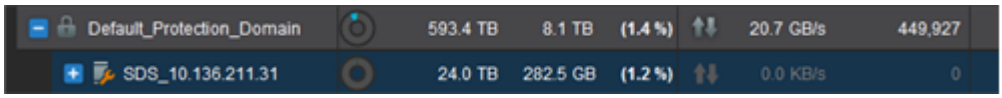

7. On an ESX node, enter maintenance mode:
   a. Log in to the vCenter via the vSphere Client or Web Client, and locate the relevant ESX IP address.
   b. Select the SVM, and from the **Basic Tasks** pane select **Shut down the virtual machine**.
   c. When the SVM is off, right-click the node and select **Enter Maintenance Mode**.
8. Obtain customer permission to reboot the node, and then gracefully 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 reboot after entering the SDS into maintenance mode.

Upgrading the firmware using the iDRAC virtual console

The iDRAC Virtual KVM console and Virtual CDROM features, provided by the iDRAC Enterprise license, eliminate the need for physical access to the ScaleIO Ready Node servers. The Hardware ISO can be attached as a remote/Virtual CDROM image and configured to perform hardware updates to the ScaleIO Ready Node firmware.

Note

For additional instructions regarding version updates following SATADOM replacement, see SATADOM firmware update on page 244.

Procedure

1. Configure a laptop with access to the iDRAC network.
2. Download the ScaleIO Ready Node Hardware Update Bootable ISO from the ScaleIO Ready Node page and make it accessible on the network share folder.
3. Connect to the iDRAC Virtual KVM console.

   For instructions, see Opening the KVM console on page 127.
4. Attach the Hardware ISO to Virtual Media.
   a. From the **Virtual Media** menu, select **Connect Virtual Media**.
   b. From the **Virtual Media** menu, select **Map CD/DVD**.
   c. Click **Browse**, and select the Hardware ISO.
5. From the **Next Boot** menu, select **Virtual CD/DVD/ISO** and configure the Next Boot option so that the server will boot to the iDRAC Virtual CDROM.

6. Power-cycle the server using the appropriate **Power** menu option, and allow the server to boot to the virtual media.
7. Repeat steps 3 on page 130 to 6 on page 131 for each ScaleIO Ready Node server that requires BIOS or firmware updates and configuration. All ScaleIO Ready Node servers can be updated in parallel.
8. When the updates are completed, refresh the iDRAC browser screen, log in to the iDRAC, and re-launch the virtual console as needed.
9. Wait for the configuration and firmware updates to complete. The server console screen will indicate when the script is complete.

   WARNING

   Do not reboot the ScaleIO Ready Node server while the update process is being performed!

   The iDRAC will be reset several times during the update process. This causes the iDRAC virtual console viewer to close, virtual media to disconnect, and the iDRAC browser window to be unavailable for several minutes during each reset. The hardware update scripts will continue to run from RAM on the server.

   The update script will generate a log indicating whether each configuration and firmware flash is successful.

10. (Optional) Check each ScaleIO Ready Node server's log for successful competition:
    a. After the update script completes, press Alt+F2 to access a user console, and then Enter to log in.
    b. Check the log contents for errors:

    ```
    less /bundleapplicationlogs/apply_components.log
    ```

    For more information, see Troubleshooting the Hardware ISO on page 248 .

    c. Press q to exit the log viewer, and then Alt+F1 to access the original console screen.
11. Reboot the ScaleIO Ready Node server and allow the update and configuration jobs to complete. Power-cycle the server using the appropriate **Power** menu option.

Return the node to operation

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

Procedure

1. Wait for the node to power on.

   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 Linux and Windows nodes, skip to the next step. On an ESX node:
   a. From the vSphere 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 **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 ESX nodes, you can migrate VMs to the node.
7. For each ScaleIO Ready Node server, after the updates are finalized, clear the iDRAC job queue using the iDRAC GUI:
   a. From your Internet browser, go to https://*<iDRAC_IP_address>*.
   b. In the **DELL Console Login** window, type these credentials:
      - username: root
      - password: calvin (for Dell-supplied nodes) or password (for EMC-supplied nodes)
   c. Click **Login**.
   d. In the iDRAC GUI navigation pane, select the **Server** node, and then select the **Job Queue** tab.
   e. Ensure that all jobs have completed successfully. Any job failures may require re-running the bootable ISO, or further troubleshooting.
   f. Select the checkbox for all of items in the **Job Queue** list, and then click **Delete**.

After you finish

Repeat for each node requiring a firmware upgrade.

CHAPTER 10

Upgrade of LSI RAID Controller Firmware and Driver

This chapter explains how to upgrade the firmware and driver on LSI RAID controllers in VxRack Node 100 Series systems. Topics include:

- LSI RAID controller upgrade on ESX 134
- LSI RAID controller upgrade on Linux 138

LSI RAID controller upgrade on ESX

Upgrade LSI RAID controller firmware and driver on ESX-based systems.

Perform the following tasks:

1. Preparation on page 134
2. Uploading firmware, driver and storcli installers on page 134
3. Upgrading the driver on page 135
4. Upgrading the firmware on page 136
5. Post-upgrade tasks on page 138

Note

This procedure applies to ScaleIO installations on ESX that are NOT using the Direct Path ability in ESX. In other words, the RAID controller driver runs on the ESX.

Preparation

Prepare your ESX-based system for the upgrade. Ensure that all relevant parties have been notified about these maintenance activities.

Before you begin

Shut down, VMotion or evacuate all VMs from the host, because it will be restarted during the upgrade procedure.

Procedure

1. SDS preparation:
   - On ScaleIO 1.32.x only, remove the SDS devices from the ScaleIO configuration (via GUI or SCLI) and wait until the rebalance is complete. (Use CTRL or SHIFT keys for multi-selection of devices in the ScaleIO GUI.) Then, power off the SVM.
   - On ScaleIO 2.x only, put the SDS into maintenance mode, and then power off the SVM.
2. Host OS preparation: put the ESX host into maintenance mode.

Uploading firmware, driver and storcli installers

Upload the RAID controller firmware, driver, and storcli installers to the host. Save them in the location `/tmp`.

Procedure

1. Download the following packages from the links listed below, unzip them, and save the required files in the location `/tmp` on your host.
   a. storcli:

      http://docs.avagotech.com/docs/1.19.04_StorCLI.zip

      Save the file `vmware-esx-storcli-1.19.04.vib` on your host.

   b. LSI_MR3 Driver:

      For ESX 6.0: https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI60-LSI-LSI-MR3-66110500-1OEM&productId=491

Save the file `lsi-mr3-6.611.05.00-1OEM.600.0.0.2768847.x86_64` on your host.

For ESX 5.5: https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI55-LSI-LSI-MR3-66110500-1OEM&productId=353

Save the file `lsi-mr3-6.611.05.00-1OEM.550.0.0.1391871.x86_64.vib` on your host.

c. LSI 3108 RAID Controller Firmware:

http://docs.avagotech.com/docs/24.15.0-0016_SAS_MR_FW_IMAGE_APP_4.650.00-6121.zip

Save the file `mr3108fw.rom` on your host.

Upgrading the driver

Upgrade the LSI RAID controller driver.

Before you begin

Ensure that you have uploaded the installer files to `/tmp` on the ESX host.

Procedure

1. Install storcli using the `vib` file from the `Vmware-NDS` folder:

```
esxcli software vib install -f -v /tmp/vmware-esx-storcli-1.19.04.vib
```

2. Install the lsi_mr3 driver:

```
esxcli software vib install -v /tmp/lsi-mr3-6.611.05.00-1OEM.600.0.0.2768847.x86_64.vib
```

3. Enable the lsi_mr3 driver:

```
esxcli system module set --enabled=true --module=lsi_mr3
```

4. If installed, disable the megaraid_sas driver:

```
esxcli system module set --enabled=false --module=megaraid_sas
```

Note

While this driver module upgrade requires a restart to take effect, it will be performed in a later step.

Upgrading the firmware

Upgrade the LSI RAID controller firmware.

Before you begin

Ensure that you have uploaded the installer files to `/tmp` on the ESX host.

The storcli executable is located in `/opt/lsi/storcli/storcli`.

Procedure

1. Pull and record the current controller configuration:

```
storcli /c0 show all
```

   Note the current firmware and driver versions shown in the output (similar to the following example):

```
Version :
=======
Firmware Package Build = 24.7.0-0026
Firmware Version = 4.270.00-3972
Bios Version = 6.22.03.0_4.16.08.00_0x060B0200
NVDATA Version = 3.1411.00-0009
Ctrl-R Version = 5.08-0006
Preboot CLI Version = 01.07-05:#%0000
Boot Block Version = 3.06.00.00-0001
Driver Name = megaraid_sas
Driver Version = 06.807.10.00-rh1
```

2. Disable CacheCade and flush all cached data to disk:

```
storcli /c0 /vall set ssdcaching=off
```

   Note

   This command will fail for SSDs. As it only applies to HDDs, this is expected behavior, and is not a problem.

3. Allow time for the cached data to write completely to disk. When this is done, VDs will show "-" in the Cac column of `storcli /c0 show` output. Once the Cac column value shows as "-" for all VDs, proceed to the next step.

   Typical output:

```
---------------------------------------------------------------
DG/VD TYPE  State Access Consist Cache Cac sCC       Size Name
---------------------------------------------------------------
1/1   RAID0 Optl  RW     Yes     RWBD  -   ON    1.090 TB VD01
2/2   RAID0 Optl  RW     Yes     RWBD  -   ON    1.090 TB VD02
3/3   RAID0 Optl  RW     Yes     RWBD  -   ON    1.090 TB VD03
```

4. Delete the CacheCade VD:

Note

Do not delete the CacheCade device until all cached data has been written to disk. Otherwise, data loss could occur.

```
storcli /c0 /v0 del cc
```

5. Enable booting with pinned cache:

```
storcli /c0 set bootwithpinnedcache=on
```

6. Upgrade the controller firmware:

```
storcli /c0 download file=/tmp/mr3108fw.rom
```

7. Restart the host to apply the firmware change:
   a. Restart the host.

   Note

   This restart will apply all of the driver and firmware upgrades.

   b. Confirm that the controller firmware and driver have been updated to the target version:

```
storcli /c0 show all | grep "Firmware Package Build\|Driver
Name\|Driver Version"
```

   Output similar to the following should be displayed:

```
Firmware Package Build = 24.15.0-0016
Driver Name = lsi-mr3 (Megaraid_sas on Linux)
Driver Version = 6.611.05.00 (6.811.02 on Linux)
```

8. Recreate the CacheCade device, then set its rdcache parameter:
   a. Type the command:

```
storcli /c0 add vd cc type=raid0 drives=E:0-1 wb
```

   b. Obtain the new VD ID in the Name column, in the row corresponding to the Cac0 type volume, using the command:

```
storcli /c0 show all
```

   This step is required because the VD ID may have changed as a result of the upgrade.

c. Type the command:

```
storcli /c0 /v<x> set rdcache=nora
```

where <x> represents the VD ID obtained in the previous step.

For example, if the VD ID in the Name column is VD1, type the following:

```
storcli /c0 add vd cc type=raid0 drives=E:0-1 wb
storcli /c0 /v1 set rdcache=nora
```

9. Activate CacheCade for the HDDs:

```
storcli /c0 /vall set ssdcaching=on
```

Post-upgrade tasks

Add SDS devices to ScaleIO.

Procedure

1. Use the vSphere Client to locate the ESX and exit Maintenance Mode.
2. Power on the SVM.
3. On ScaleIO v1.32.x systems only, using the GUI or SCLI, add the SDS devices to ScaleIO.
4. On ScaleIO v2.x systems only, use the GUI **Backend** view to locate the SDS. Right-click the SDS, and select **Exit Maintenance Mode**.

After you finish

Wait until rebuild and rebalance activities are complete before proceeding to the next node.

LSI RAID controller upgrade on Linux

Upgrade LSI RAID controller firmware and driver on Linux-based systems.

Perform the following tasks:

1. Preparation on page 138
2. Uploading firmware, driver and storcli installers on page 139
3. Upgrading the driver on page 139
4. Upgrading the firmware on page 140
5. Post-upgrade tasks on page 143

Preparation

Prepare your Linux-based system for the upgrade. Ensure that all relevant parties have been notified about these maintenance activities.

Procedure

1. SDS preparation:

- On ScaleIO 1.32.x only, remove the SDS devices from the ScaleIO configuration (via GUI or SCLI) and wait until the rebalance is complete. (Use CTRL or SHIFT keys for multi-selection of devices in the ScaleIO GUI.)
- On ScaleIO 2.x only, put the SDS into maintenance mode.

2. Host OS preparation: shut down services and applications as needed, in preparation for maintenance.

Uploading firmware, driver and storcli installers

Upload the RAID controller firmware, driver, and storcli installers to the host. Save them in the location `/tmp`.

Procedure

1. Download the following packages from the links listed below, unzip them, and save the required files in the location `/tmp` on your host.

   a. storcli:

   http://docs.avagotech.com/docs/1.19.04_StorCLI.zip

   Save the file `storcli-1.19.04-1.noarch.rpm` on your host.

   b. Megaraid_sas driver:

   http://docs.avagotech.com/docs/MR_LINUX_DRIVER_6.11-06.811.02.00-2

   Save the file `kmod-megaraid_sas-06.811.02.00_xyz.x86_64.rpm` on your host (distribution/release specific; see installer archive).

   c. LSI 3108 RAID Controller Firmware:

   http://docs.avagotech.com/docs/24.15.0-0016_SAS_MR_FW_IMAGE_APP_4.650.00-6121.zip

   Save the file `mr3108fw.rom` on your host.

Upgrading the driver

Upgrade the LSI RAID controller driver.

Before you begin

Ensure that you have uploaded the installer files to `/tmp` on the Linux host.

If the SDS is running ScaleIO v2.x, ensure that it is in maintenance mode.

Procedure

1. Upgrade the megaraid_sas 6.811.02 driver:

```
rpm -Uvh /var/tmp/kmod-
megaraid_sas-06.811.02.00_el7.2-2.x86_64.rpm
```

Note

This installs the driver to the base kernel version for the relevant distribution/release.

2. Stop the SDS:

```
/opt/emc/scaleio/sds/bin/delete_service.sh
```

3. Unload the driver:

```
modprobe -r megaraid_sas
```

4. Copy the driver that you just installed to the base kernel version folder for your distribution/release, and overwrite the current driver file:

```
cp
/lib/modules/3.10.0-327.el7.x86_64/extra/megaraid_sas/
megaraid_s
as.ko /lib/modules/`uname
-r`/kernel/drivers/scsi/megaraid/megaraid_sas.ko
```

```
cp: overwrite
"/lib/modules/3.10.0-327.4.5.el7.x86_64/kernel/drivers/scsi/
mega
raid/megaraid_sas.ko"?
Y
```

5. Load the driver:

```
modprobe megaraid_sas
```

6. Confirm that the just-installed driver is loaded:

```
modinfo megaraid_sas |grep version
```

```
version: 06.811.02.00
```

Upgrading the firmware

Upgrade the LSI RAID controller firmware.

Before you begin

Ensure that you have uploaded the installer files to `/tmp` on the Linux host.

The storcli executable is located in `/opt/MegaRAID/storcli/storcli64`.

Procedure

1. Pull and record the current controller configuration:

```
storcli64 /c0 show all
```

Note the current firmware and driver versions shown in the output (similar to the following example):

```
Version :
=======
Firmware Package Build = 24.7.0-0026
Firmware Version = 4.270.00-3972
Bios Version = 6.22.03.0_4.16.08.00_0x060B0200
NVDATA Version = 3.1411.00-0009
Ctrl-R Version = 5.08-0006
Preboot CLI Version = 01.07-05:#%0000
Boot Block Version = 3.06.00.00-0001
Driver Name = megaraid_sas
Driver Version = 06.807.10.00-rh1
```

2. Disable CacheCade and flush all cached data to disk:

```
storcli64 /c0 /vall set ssdcaching=off
```

Note

This command will fail for SSDs. As it only applies to HDDs, this is expected behavior, and is not a problem.

3. Allow time for the cached data to write completely to disk. When this is done, VDs will show "-" in the Cac column of `storcli64 /c0 show` output. Once the Cac column value shows as "-" for all VDs, proceed to the next step.

   Typical output:

```
-------------------------------------------------------------
DG/VD TYPE  State Access Consist Cache Cac sCC       Size Name
-------------------------------------------------------------
1/1   RAID0 Optl  RW     Yes     RWBD  -   ON    1.090 TB VD01
2/2   RAID0 Optl  RW     Yes     RWBD  -   ON    1.090 TB VD02
3/3   RAID0 Optl  RW     Yes     RWBD  -   ON    1.090 TB VD03
```

4. Delete the CacheCade VD:

Note

Do not delete the CacheCade device until all cached data has been written to disk. Otherwise, data loss could occur.

```
storcli64 /c0 /v0 del cc
```

5. Enable booting with pinned cache:

```
storcli64 /c0 set bootwithpinnedcache=on
```

6. Upgrade the controller firmware:

```
storcli64 /c0 download file=/tmp/mr3108fw.rom
```

7. Restart the host to apply the firmware change:
   a. Restart the host.

      Note

      This restart will apply all of the driver and firmware upgrades.

   b. Confirm that the controller firmware and driver have been updated to the target version:

```
storcli64 /c0 show all | grep "Firmware Package Build\|
Driver
Name\|Driver Version"
```

      Output similar to the following should be displayed:

```
Firmware Package Build = 24.15.0-0016
Driver Name = lsi-mr3 (Megaraid_sas on Linux)
Driver Version = 6.611.05.00 (6.811.02 on Linux)
```

8. Recreate the CacheCade device, then set its rdcache parameter:
   a. Type the command:

```
storcli64 /c0 add vd cc type=raid0 drives=E:0-1 wb
```

   b. Obtain the new VD ID in the Name column, in the row corresponding to the Cac0 type volume, using the command:

```
storcli64 /c0 show all
```

      This step is required because the VD ID may have changed as a result of the upgrade.

   c. Type the command:

```
storcli64 /c0 /v<x> set rdcache=nora
```

      where <x> represents the VD ID obtained in the previous step.

      For example, if the VD ID in the Name column is VD1, type the following:

```
storcli64 /c0 add vd cc type=raid0 drives=E:0-1 wb
storcli /c0 /v1 set rdcache=nora
```

9. Activate CacheCade for the HDDs:

```
storcli64 /c0 /vall set ssdcaching=on
```

Post-upgrade tasks

Post-upgrade tasks for the SDS.

Procedure

1. Start the SDS:

```
/opt/emc/scaleio/sds/bin/create_service.sh
```

2. If the SDS is running ScaleIO v2.x , exit maintenance mode.
3. If the SDS is running ScaleIO v1.32.x, use either the GUI or SCLI to add the ScaleIO devices back to the SDS.
4. Wait until rebuild and rebalance activities are complete before proceeding to the next node.

PART 5

Reference Topics

This part describes various software and hardware upgrade tasks, many of which are referred to in other sections of this document. Topics include:

Chapter 11, "Post-Upgrade and Other Related Activities"

Chapter 12, "Troubleshooting"

Chapter 13, "Frequently Asked Questions"

Chapter 14, "DTK - Hardware Update Bootable ISO"

CHAPTER 11

Post-Upgrade and Other Related Activities

This chapter describes various activities that are related to upgrading. Topics include:

- System analysis overview 148
- System requirements 155
- Extracting ScaleIO packages 160
- Upgrading the ScaleIO GUI 160
- Upgrading the SDC 160
- Manually upgrading the SDC version in ESX environment 161
- Modifying SDC parameters 162
- Non-disruptive SNMP upgrade issues 163
- Installing with the full Installation Manager 164
- Configuring Installation Manager properties 170
- Upgrading the Gateway when a custom certificate is used 171
- Upgrading the BMC firmware on a VxRack Node 100 Series node 171
- Unregistering the ScaleIO plug-in 172
- Registering the ScaleIO plug-in manually 172
- Installing the SDC on ESX hosts 174
- Upgrading Windows servers when the OS is installed on SATADOM 176
- Troubleshooting plug-in registration issues 176
- SVM manual memory allocation 177
- Manual migration of an RDM-based ScaleIO system to DirectPath-based 178
- Deploying on Ubuntu, Oracle Linux, or CoreOS servers 180
- Maintaining a ScaleIO system 186
- Configure ESRS after upgrading 198

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 only RHEL 6.x servers, with the following 3rd-party tools:
  - Netcat
  - StorCLI
  - PERCCLI
  - smartctl
- Requires a ScaleIO Gateway server:
  - On a Linux or Windows server. ScaleIO is tested on RHEL 6.x and RHEL 7.x as well as Windows2K12 servers.
  - With 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 network configuration.

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 IPv4 and IPv6, the IPv6 (only) analysis is not conclusive; IPv6 is not supported.

- 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

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 Home** 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

   c. Click **Retrieve system topology**.

      The system topology is displayed.

d. Click **Analyze System**.

The system analysis begins.

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.

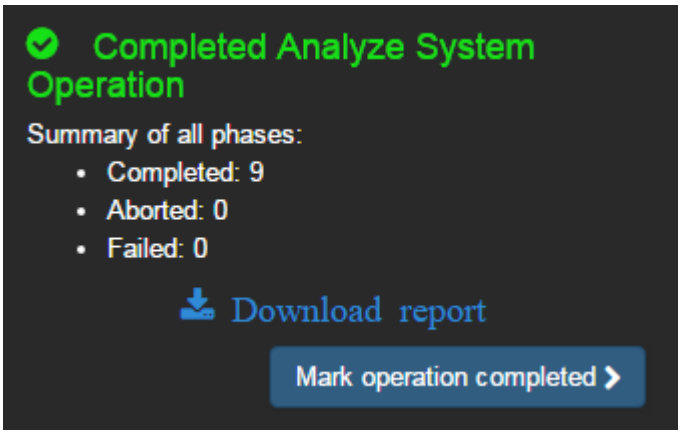

The report is saved, in ZIP format, in the default download location of the server on which the report was run.

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.

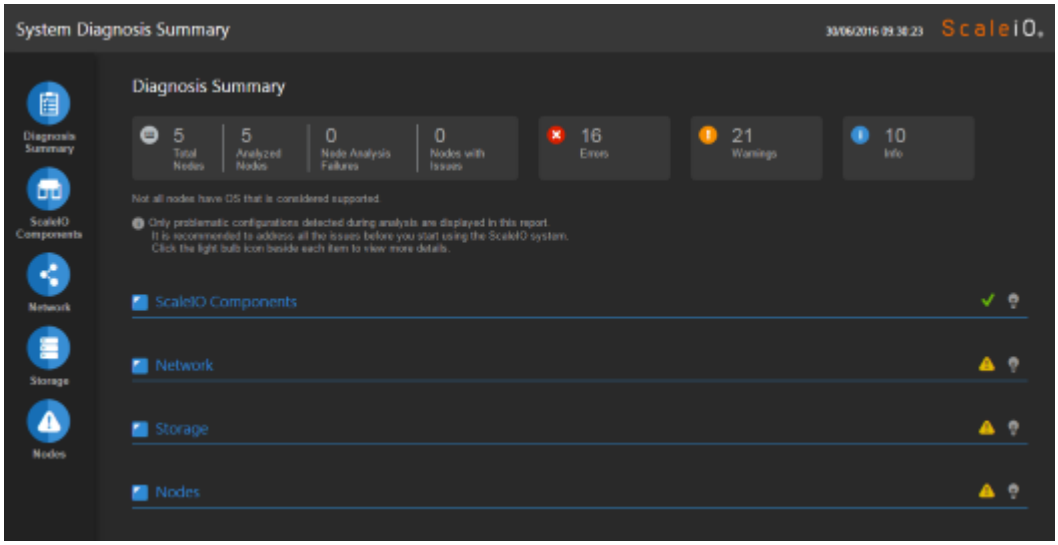

System analysis report description

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

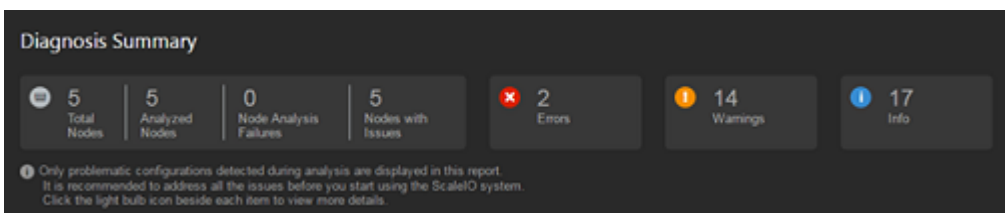

When the analysis first opens, the major sections are shown in summary form. You can expand them to show detailed diagnostic reports, as follows:

- 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.

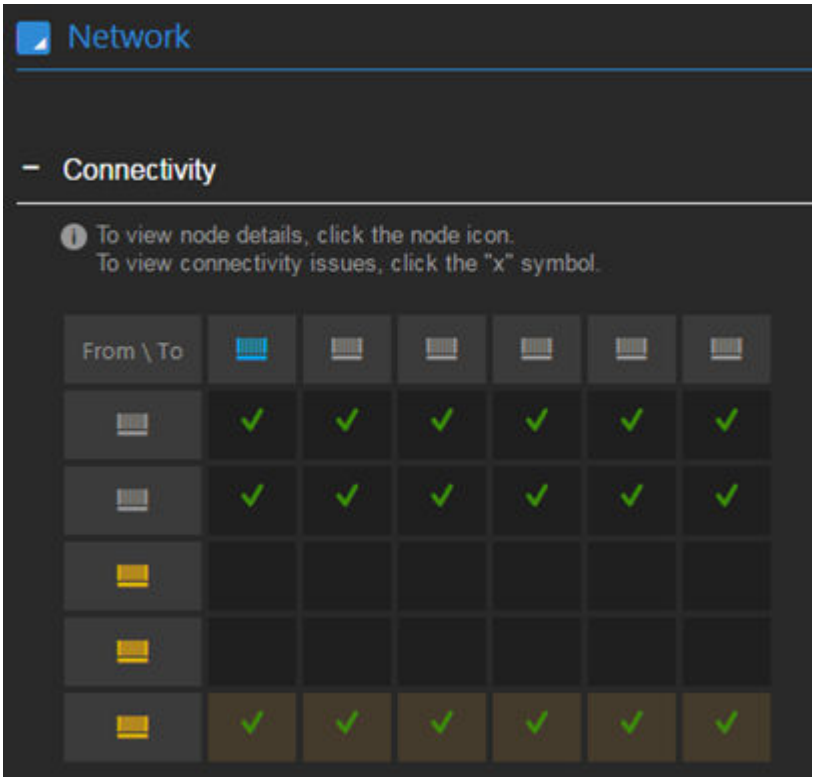

  - Connectivity
    Performing pings between ScaleIO components leads to detecting and resolving connectivity-related issues in the system. If the regular pings succeed, then the MTU pings, followed by the Ncat pings are performed to the ports used by the ScaleIO application.

    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.

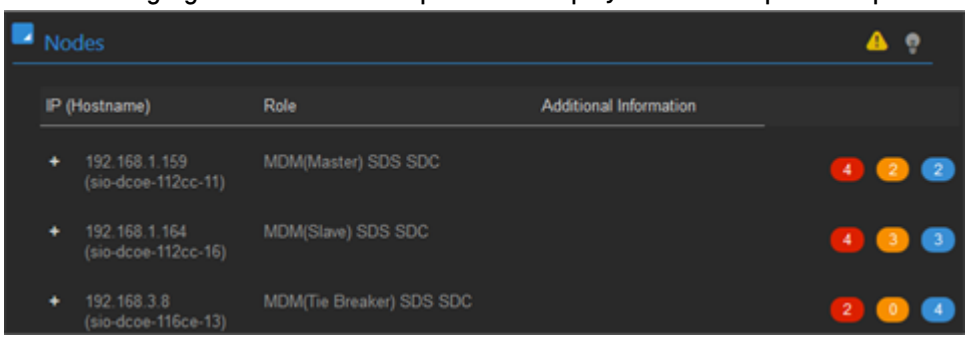

You can show additional details of an error by clicking the ( ) icon. A pop-up window similar to the following is displayed.

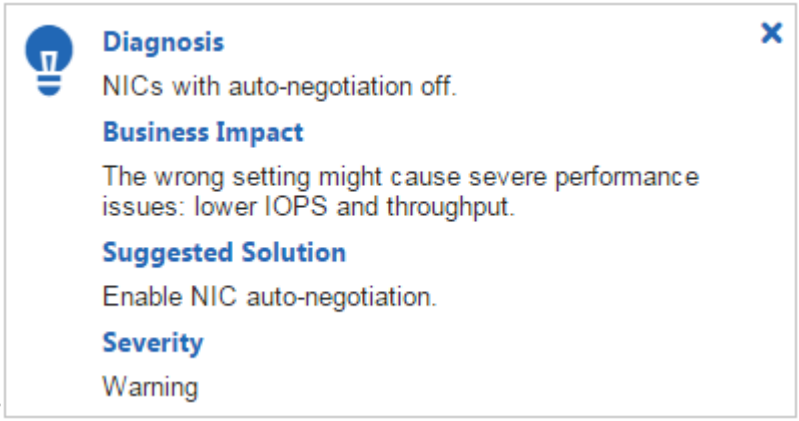

The following table describes how to navigate the analysis information:

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
or	Connectivity issues were detected (for more information, click the symbol)
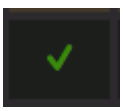 (black background)	Connectivity matrix background: all pertinent connectivity tools are available
(brown background)	Connectivity matrix background: connectivity tests could not be performed, due to a missing tool
	Error counters: • Red - Error • Orange - Warning

Symbol / Interface Elements	Description
	• 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.

System requirements

This section lists the requirements for system components.

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.

Physical server requirements

Table 5 Server physical requirements

Component	Requirement
Processor	One of the following: • Intel or AMD x86 64-bit (recommended) • Intel or AMD x86 32-bit (for Xen only)
Physical memory	ScaleIO component requirements: • 500 MB RAM for the Meta Data Manager (MDM) • 500 MB RAM for each ScaleIO Data Server (SDS) • 50 MB RAM for each ScaleIO Data Client (SDC) DAS Cache memory requirements (ScaleIO Ready Node, non-XenServers only). Add to every SVM/node that will be using DAS Cache: • 1U1N servers—500 MB RAM • 2U1N servers—1 GB RAM To calculate SVM memory allocation, use the formulas provided in the *ScaleIO Deployment Guide*.
Disk space	• 1 GB for each physical node or Xen hypervisor • 10 GB for VMware topologies

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

Table 5 Server physical requirements (continued)

Component	Requirement
	▪ SNMP traps: 162 **Note** You can change the default ports. For more information, see "Changing default ports" in the user documentation.

Supported operating systems

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

Table 6 Supported operating systems - ScaleIO components

Operating system	Requirement
Linux	Supported versions: • CentOS 6.x-7.x, Oracle Linux 6.5/7.x • Red Hat 6.x-7.x • SLES 11.3, SLES 11.4, SLES 12, SLES 12.1 • Ubuntu 14.04, Ubuntu 16.04 **Note** Before deploying on Ubuntu servers, you must prepare the environment, as described in the *EMC ScaleIO Deployment Guide*. Packages required for all components, all Linux flavors: • numactl • libaio Additional packages required for MDM components: • bash-completion (for SCLI completion) • Latest version of Python 2.X When installing the MDM component on Linux CentOS 6 or RHEL 6 hosts (for software-only systems), set the shared memory parameter in the /etc/sysctl.conf file to at least the following value: kernel.shmmax=209715200. To use this value, type the `sysctl -p` command. To use the secure authentication mode, ensure that OpenSSL 64-bit v1.0.1 or later is installed on all servers in the system. To use the secure authentication mode on SLES 11.3 or SLES 11.4 servers, ensure that the OpenSSL on the server is v1.0.1 or greater, or install these packages (from the ISO in the Complete VMware SW download container) on the server: • `libopenssl1_0_0-1.0.1g-0.40.1.x86_64.rpm` • `openssl1-1.0.1g-0.40.1.x86_64.rpm` To use LDAP, ensure that OpenLDAP 2.4 is installed on all servers.

Table 6 Supported operating systems - ScaleIO components (continued)

Operating system	Requirement
Windows	Supported versions: • 2008 R2, 2012, 2012 R2, or 2016 (in v2.0.1.1 and later). Server Core editions are not supported. • For VxRack Node 100 Series, only 2012 R2 is supported. • On all MDM servers, install the EMC-provided `PythonModulesInstall.exe` on all MDM nodes. The file is supplied on the ISO, or download from the EMC Online Support site (search for ScaleIO Python Installation Modules) on https://support.emc.com. • To install SDC or RFcache on 2008 R2, ensure that Microsoft Security Update KB3033929 is installed. To use the secure authentication mode, ensure that these are installed on all servers in the system: • OpenSSL 64-bit v1.0.1 or later • Visual C++ redistributable 2010 package, 64-bit To use RFcache, ensure that Visual C++ redistributable 2010 package, 64-bit is installed on all servers in the MDM cluster and on all SDSs.
Hypervisors	• VMware ESXi OS: 5.5 or 6.0, managed by vCenter 5.5 (update 2e or later) or 6.0 • Hyper-V • Xenserver 6.5 or 7.0 **Note** OpenSSL 64-bit v1.0.1 is supported on XEN 6.5 SP1 (or later) • RedHat KVM
OpenStack	Supported versions: • Liberty • Mitaka • Newton • Redhat OpenStack Platform • Mirantis OpenStack • Canonical Ubuntu OpenStack.

GUI server requirements

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

Component	Requirement
	• Screen resolution: 1366 x 768 minimum

ScaleIO Gateway server requirements

Component	Requirement
Supported operating systems	• Windows 2008 R2, 2012 R2, or 2016, including the Visual C++ redistributable 2010 package, 64-bit • Linux: ▪ CentOS 6.x-7.x ▪ Oracle Linux 6.5/7.x ▪ Red Hat 6.x-7.x ▪ SLES 11.3, SLES 12, SLES 12.1 ▪ Ubuntu 14.04, Ubuntu 16.04 Every server requires 2 cores and a minimum of 3 GB available RAM.
Supported web browsers	• Internet Explorer 10, or later • Firefox, version 42, or later • Chrome, version 45, or later
Java requirements	• 1.8 or later, 64-bit
Other	• For a Windows Installation Manager (IM): ▪ The WMI service must be enabled on the IM server and on all Windows ScaleIO nodes. ▪ Don't install the Gateway on a server on which RFcache (the `xcache` package) or 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*.

Extracting ScaleIO packages

Extract ScaleIO Ubuntu (DEB) or CoreOS (BSX) packages.

ScaleIO packages for Ubuntu and CoreOS must be extracted for some installation and upgrade scenarios.

Procedure

1. Extract a package:

OS	Description

CoreOS	a. Untar the package:
	`tar -xvf EMC-ScaleIO-<component>-2.0-13000.X.CoreOS.x86_64.tar`
	b. Extract the package to yield the BSX file:
	`./siob_extract EMC-ScaleIO-<component>-2.0-13000.X.CoreOS.x86_64.siob`

Ubuntu	a. Untar the package:
	`tar -xvf EMC-ScaleIO-<component>-2.0-13000.X.Ubuntu.x86_64.tar`
	b. Extract the package to yield the DEB file:
	`./siob_extract EMC-ScaleIO-<component>-2.0-13000.X.Ubuntu.x86_64.siob`

2. Repeat the previous step for all packages that need to be extracted.

Upgrading the ScaleIO GUI

You can upgrade the ScaleIO GUI.

Procedure

1. Upgrade the GUI:

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

Upgrading the SDC

You can use the VMware plug-in to upgrade an SDC that is installed directly on an ESXi server.

Note

If the currently installed ScaleIO system includes any SDC nodes that are installed in a ScaleIO virtual machine (SVM), you must switch those SDCs to an ESX before upgrading the SDC. To switch SDCs, contact EMC Support.

To upgrade the SDC version, perform the following:

Procedure

1. From the VMware plug-in **SDCs** screen, right-click an SDC and select **Upgrade SDC**.

   The **Upgrade SDC** dialog box appears.

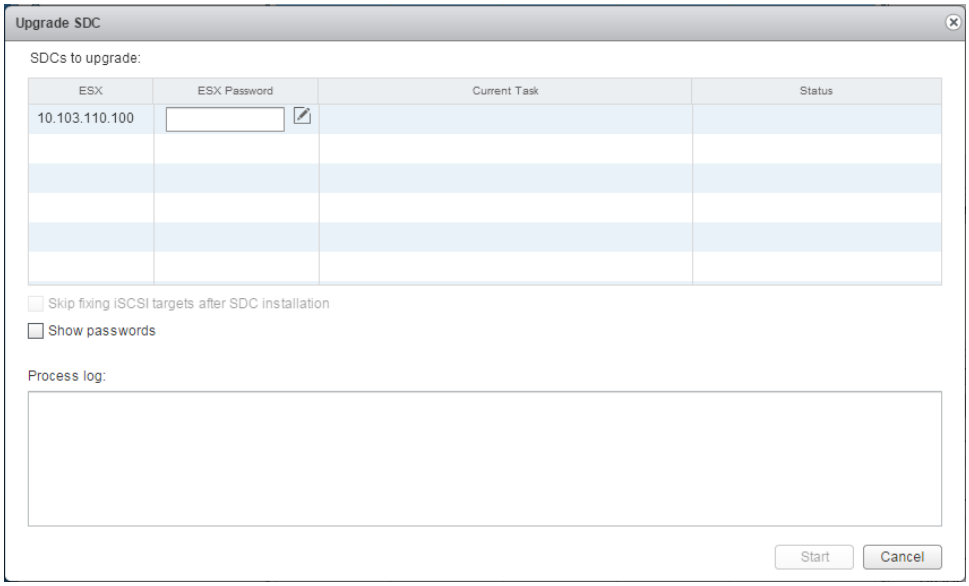

2. Enter the ESX root password, then click **Start**.
3. If prompted, select the new SDC driver, then click **OK**.
   a. To upgrade on an ESXi 5.5 server, select `sdc-2.0-13000.X-esx5.5.zip`
   b. To upgrade on an ESXi 6.0 server, select `sdc-2.0-13000.X-esx6.0.zip`

   Upgrade progress is shown in the **Upgrade SDC** screen. You can allow this process to run in the background. To check on progress, click **Show SDC upgrade process** from the main ScaleIO screen.

Manually upgrading the SDC version in ESX environment

Upgrade the SDC on ESX manually.

Procedure

1. Log in to the ESX and set the acceptance level:

```
esxcli software acceptance set-level=PartnerSupported
```

2. If your ScaleIO system supports SDS maintenance mode (v2.0 and later), log in to the SDS and place the SDS in maintenance mode, with the CLI, REST, or GUI. If not, skip this step.
3. Power off the ScaleIO virtual machine (SVM).
4. Place the ESX host in maintenance mode.

   This operation migrates all the guest VMs.
5. Install and upgrade the VIB:

```
esxcli software vib update -d "Full Path"
```

6. Restart the ESX host.

   Note

   This step is critical for implementing the SDC version upgrade.

7. Using the vSphere client, remove the ESX host from maintenance mode.
8. Using the vSphere client or the GUI, remove the SDS from maintenance mode.

Results

The SDC version is upgraded.

Modifying SDC parameters

Modify SDC parameters, for example, the virtual IP address of an MDM cluster.

The modification causes an SDC to use a different MDM virtual IP address or to update the SDC performance settings and configuration.

Procedure

1. Copy the drv_cfg tool to the node.

   The drv_cfg tool is not a part of the artifacts available on the support site. Contact https://support.emc.com to access this tool and the guidelines for how to use it.
2. If needed, update the SDC performance settings and the existing configuration.

   This step is needed only if you change the settings after an upgrade.
3. View the complete list of SDC performance settings for the scini driver:

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

4. Update the scini settings:

```
esxcli system module parameters set -m scini -p "...."
```

CAUTION

You must specify all parameters in this command when using the -p option. Any parameter that existed, that is not specified now will be replaced with 0 or an empty string.

Example for modifying virtual IP address:

```
esxcli system module parameters set -m scini -p
"IoctlIniGuidStr=7ea77fc9-6125-44c1-9746-9d20053aed93
IoctlMdmIPStr=9.10.21.141,9.10.21.142,9.10.121.142,9.10.121.14
1"
```

5. You can us the following cat command to get the full string output. This enables you to add or modify only the parameters you need.

```
cat /etc/vmware/esx.conf |grep scini/
 /vmkernel/module/scini/options =
"IoctlIniGuidStr=7ea77fc9-6125-44c1-9746-9d20053aed93
IoctlMdmIPStr=9.201.112.254,9.201.111.254"
```

Copy/paste the info between the quotation marks (" ") and modify; otherwise, it will return additional values.

Results

Selected SDC parameters are modified.

Non-disruptive SNMP upgrade issues

The SNMP trap sender's credentials are handled differently in ScaleIO v1.32 and v2.0. This needs to be taken into account during or after the upgrade process.

If a version 1.32 ScaleIO Gateway is configured as an SNMP trap sender, it is configured with SNMP MDM credentials. When upgrading to version 2.0, the SNMP MDM credentials will be uploaded to the lockbox, and will be erased from the old gatewayUser.properties configuration file. The update is performed using the SNMP MDM password as a phrase, which may be considered as not meeting the ultimate security guidelines (the same text is used for both the MDM password value and the phrase for Lockbox secrets). However, it is a value which should be confidential, and is known to the user.

To harden the password, one of the following procedures can be performed:

- After upgrading the ScaleIO Gateway, change the Lockbox phrase using SioGWTool
- During the ScaleIO Gateway upgrade, provide the LOCKBOXPHRASE env variable as part of the upgrade command.

Examples:

- *Windows:* >> **msiexec /i EMC-ScaleIO-gateway-2.0-4000.0-x64.msi /L*v GW_install.log LOCKBOXPHRASE=user_phrase**

- *Linux:* `>>LOCKBOXPHRASE=user_phrase rpm -Uvh EMC-ScaleIO-gateway-2.0-4000.0.x86_64.rpm`

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 *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.

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:

- "Log in to the IM server"
- "Upload installation packages"
- "Upload CSV file"
- "Initiate the installation"
- "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/elnhome.

   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 *ScaleIO Deployment Guide*.
Configuration options	Enable zero-padding	Enable zero-padding of new Storage Pools.

Option type	Option	Description
		Zero padding is required for using the background scanner in data comparison mode and for use with RecoverPoint splitter replication. **Note** Replication support is version-specific; see the ESSM for full details.
	Enable alert service	Enable the alert service required for SNMP and ESRS reporting. This also creates and configures the lockbox. If you select this option, the **Traps Receiver IP/ 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	Disable the need for secure communication mode between management clients and the MDM. **Note** Disabling secure communication has security implications, described in "Using SCLI in non-secure mode".
	Disable secure communication with LIA	Disable the need for secure communication mode between management clients and the LIA. **Note** Disabling secure communication has security implications, described in "Using SCLI in non-secure mode".
	Disable authentication in internal components	Disable the need for authentication between SDSs and MDMs.

Option type	Option	Description
	Use trusted IP addresses	Limit the ScaleIO Gateways that can communicate with this LIA. 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.

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)
6. 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**.

7. Click the **Monitor** tab.

Complete the installation

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

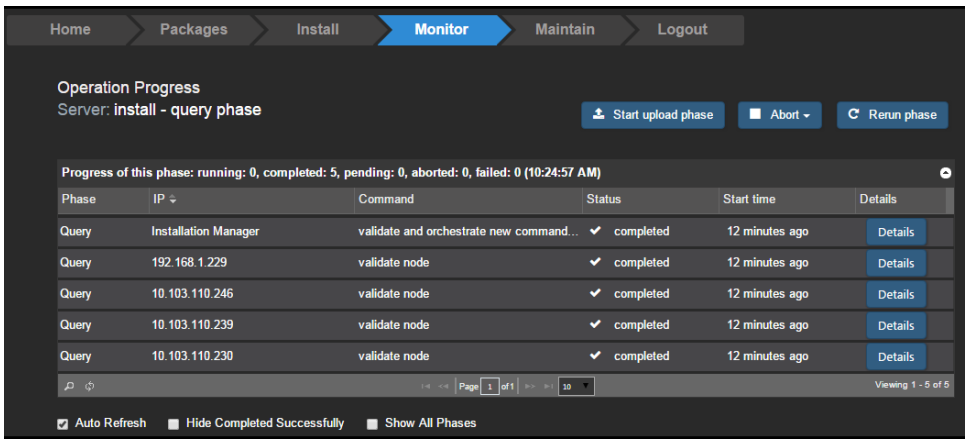

In the query stage, the IM validates that there are no previously installed ScaleIO components on any of the requested installation nodes.

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.

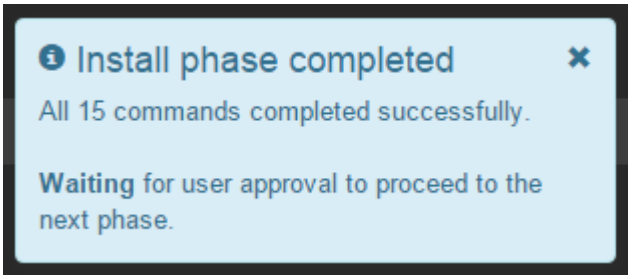

If an error message is displayed during the process, you can retry the failed operations (with the **Retry failed** button), continue to the next phase, or abort the installation.

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 analyis overview."

Best practice suggestion

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

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 ScaleIO user documentation.

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

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

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

```
service scaleio-gateway restart
```

Results

Configuration is complete.

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`

Upgrading the BMC firmware on a VxRack Node 100 Series node

To upgrade the BMC firmware on a VxRack Node 100 Series node, you are required to perform the following steps:

Before you begin

- Ensure that the device performing the upgrade resides in the same LAN as the BMC that you are upgrading.
- Make sure that you have the following information at hand:
  - The IP address of the BMC port
  - The username and password for accessing the BMC

Procedure

1. Download the `2b2s2344.ima_enc` file to the device that will perform the upgrade. The file is located at https://support.emc.com/downloads/39045_VxRack-Node/.
2. From your Internet browser, go to http://*<BMC_Port_IP_address>*.
3. In the **Quanta Console Login** window, log in using the BMC username and password.
4. On the **Dashboard** tab, check the firmware revision.

   If the dashboard displays firmware revision 23.44.00, which is the newest version, you do not need to perform an upgrade.
5. On the **Maintenance** tab, click **BMC Firmware Update**.
6. Ensure that the **Preserve all Configuration** check box is not selected, and then click **Enter upgrade mode**.
7. Click **OK**.

   The **BMC Firmware Update window** appears.
8. When the **Upload Firmward** message appears, click **Browse**.
9. Browse to the `2b2s2344.ima_enc` file stored locally on your device, select the file, and then click **Upload**.
10. At the prompt, click **Proceed**.
11. Click **OK** to start the upgrade operation.
12. Wait for the upgrade to finish.

    A reset will be performed at the end of the operation.
13. Wait 3 minutes, and then open a new internet connection to the BMC.
14. Log in and verify on the dashboard that the new version is 23.44.00.

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.

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-2.0-13000.X.zip`
5. Upload `EMC-ScaleIO-vSphere-web-plugin-2.0-13000.X.zip` to the HTTP/HTTPS server.
6. From PowerCLI, run the following script: `ScaleIOPluginSetup-2.0-130000.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
```

Installing the SDC on ESX hosts

This topic describes how to install the SDC component on ESX hosts. This is a prerequisite for deploying a new ScaleIO system, and it must be applied to every ESX host that will be in the ScaleIO system, regardless of what role it is playing.

This process installs the SDC on the ESX host, after which you must restart the ESX host.

To install the SDC component, perform the following:

Procedure

1. From the **Basic tasks** section of the **EMC** ScaleIO screen, click **Install SDC on ESX**.

   The **Install SDC on ESX** screen appears, showing all the ESX hosts on all the connected vCenters:

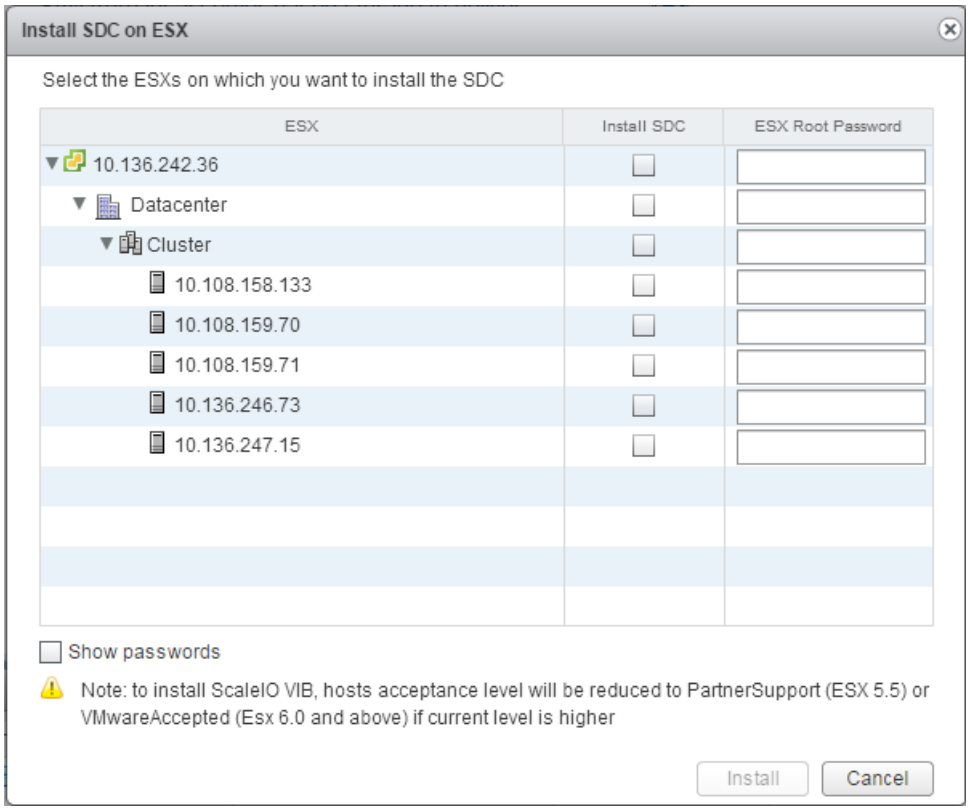

Note

Entering the cluster password will automatically enter the same password for all servers in the cluster.

2. Select the ESX hosts, and enter the root password for each host.

Note

It is highly recommended to select all ESX hosts that may be included in an ESX system, even if only in the future.

3. Click **Install**.

   The status appears in the dialog.

4. When finished, click **Finish**.
5. Restart each ESX host.

Note

You must restart the ESX hosts before proceeding.

6. 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**.

After you finish

Proceed with the ScaleIO deployment.

Upgrading Windows servers when the OS is installed on SATADOM

Procedures to enable upgrading ScaleIO components from v2.0.0.2 (or later) , when the Windows OS is installed on the SATADOM.

You must perform these steps before upgrading the ScaleIO components.

Procedure

1. From the Windows Start menu, select **System**.
2. Click **Advanced system settings**.
3. In the **Performance** section, click **Settings**.
4. From the **Advanced** tab, in the **Virtual memory** section, click **Change**.
5. Clear the **Automatically manage paging file size for all drives** option.
6. Set a custom size of 4096, for both the initial and maximum size.
7. Click **OK** to approve changes.

Results

You can now begin the upgrade process.

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

vCenter	Operating system	Path to file
	Linux	/opt/.vmware/scaleio
6.x	Windows	C:\Users\vspherewebclientsvc \AppData\Roaming\VMware\scaleio
	Linux	/etc/vmware/vsphere-client/vc-packages/scaleio

SVM manual memory allocation

When using the plug-in for a clean deployment, SVM memory allocation is performed automatically. In the following cases, SVM memory allocation must be performed manually:

- Manual deployment on VMware.
- Extending an existing SVM with a new ScaleIO role/component, whether this is being done with the plug-in or manually.
  Workaround: Perform all the parts of step 1 and step 2 before extending the additional role/component on the SVM. Perform the steps on one SVM at a time.
- Changing the SDS performance profile, post deployment, or after an upgrade from v1.32.x to v2.0.
  Workaround: Perform all the parts of step 1 one SVM at a time.
- Post backend upgrade from v1.32.x to v2.0.
  Workaround: Perform all the parts of step 1 and step 2 on all SVMs that were upgraded, in the following order, one SVM at a time.
- Adding capacity to an SDS that was deployed with partially-populated capacity.
  Workaround: Perform all the parts of step 1 one SVM at a time.

Procedure

1. For SVMs that are SDS-only, perform the following:
   a. Move the SDS to maintenance mode (MM).
   b. Shut down the SVM.
   c. Increase SVM memory, according to the formula below.
   d. Power up the SVM.
   e. Exit MM.
2. For SVMs that are MDM (Master, Slave, or TB, may contain SDS, also):
   a. Start with Slaves and TBs:
      a. Move the SDS to maintenance mode (MM).
      b. Shut down the SVM.
      c. Increase SVM memory, according to the formula below.
      d. Power up the SVM.
      e. Exit MM.
   b. Proceed with the Master MDM:
      a. Switch ownership, so the Master MDM is now a Slave MDM.

b. Move the SDS to maintenance mode (MM).
c. Shut down the SVM.
d. Increase SVM memory, according to the formula below.
e. Power up the SVM.
f. Exit MM.

The memory allocation formula:

Component	Memory allocation rules		
Base SVM	• 350 MB		
MDM (Master/Slave)	• 470 MB + (500 KB * 8 TB of volume capacity) + (1.44 KB * number of volumes) + (4 KB * number of SDS devices) • Maximum supported volumes: 256 K		
Tie Breaker MDM	• 50 MB		
SDS	• (Base) 536 MB + (RmCache Size) * 1.15 + (Storage capacity in TB) * 53 MB • For SDS high performance profile, we add 195 MB.		
SDC	• 132 KB + 23 MB * (number of MDMs) + 25 KB * (number of SDSs) + 1.5 KB * (number of volumes) + 16 B * (number of volume blocks) + 24 KB * (8 TB of volume capacity) • Volume blocks: 1 GB storage = 8 volume blocks		
RFcache	• 16 * (cache_size/page_size) • Commonly-used sizes:		
	RFcache page size	RFcache memory requirement, if the cache device is 800 GB	RFcache memory requirement, if the cache device is 1.6 TB
	64 K	200 MB	400 MB
	32 K	400 MB	800 MB
	16K	800 MB	1.6 GB
	8 K	1.6 GB	3.2 GB
	4 K	3.2 GB	6.4 GB

Manual migration of an RDM-based ScaleIO system to DirectPath-based

On PowerEdge servers, use this procedure to manually migrate a ScaleIO v2.0.1.4 (or later) system from RDM-based device management to DirectPath-based

(passthrough) device management. IO is not interrupted when following the instructions provided.

Before you begin

- The current ScaleIO system must be v2.0.1.4 or later with RDM-based device management. If you are migrating systems that are earlier than v2.0.1.4, you must install the RAID controller tools before adding devices, as described in "Install RAID controller tools."
- You need to be able to access the vCenter.
- You need to be able to access the ScaleIO management clients.
- It is recommended to schedule this procedure during a maintenance window.

This manual procedure must be performed on every SDS node in the system. First, you enable the passthrough mode on the ESX node, then add PCI devices to the ScaleIO VM (SVM) on that server. Then, repeat this procedure for the next SDS node.

Procedure

1. Detach RDM devices from the SVM:
   a. From ScaleIO, use the GUI (or CLI) to select any SDS in the system.
   b. Enter the SDS into maintenance mode.
   c. From the vCenter, power-off the SVM.
   d. From the vCenter, edit the SVM and remove all the attached RDM disk devices.

      Use the remove - not the remove and delete - option.
2. Enable passthrough on the ESX node:
   a. From the vCenter **Configuration** > **Advanced settings** view, select **Enable passthrough**.
   b. Select the controller (RAID, SAS, or HBA) that is managing the disks used by ScaleIO.
3. Configure the BIOS of the ESX node for passthrough mode:
   a. In the vCenter Host view, enter the node into maintenance mode.
   b. Once maintenance mode has started, reboot the node.

   The selected ESX node is configured for DirectPath/passthrough mode.
4. Add PCI devices to the SVM:
   a. After the ESX node reboot completes, remove the node from maintenance mode.
   b. Edit the SVM, and add PCI devices.
   c. Power on the SVM.

      Added PCI devices are now displayed in the SVM.
   d. From ScaleIO, remove the SDS from maintenance mode.
   e. Wait for rebuild and rebalance activities to complete.
5. After all rebuild and rebalance activities are finished on one SDS, repeat this entire procedure on every SDS, one-at-a-time.

After you finish

If you are migrating systems that are earlier than v2.0.1.4, you must install the RAID controller tools before adding devices, as described in "Install RAID controller tools." You can now use the ScaleIO management clients to add more devices to your SDSs. For more information, see the *ScaleIO User Guide*.

Deploying on Ubuntu, Oracle Linux, or CoreOS servers

This section describes steps that must be taken when deploying SDC or RFcache on Ubuntu, Oracle Linux (OL), or CoreOS servers. You must ensure that the ScaleIO SDC and RFcache (XC) drivers are compatible with the Ubuntu/OL/CoreOS kernel.

Note

All references to Ubuntu in this section apply equally to supported versions of OL and CoreOS.

To ensure that a client node always has an SDC/XC kernel driver that matches the client kernel, the SDC/XC startup script uses `driver_sync.sh`, a synchronization utility that fetches SDC/XC 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/XC 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/XC 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/xcache 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://QNzgdxXix: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://
QNzgdxXix: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 7 driver_sync.conf parameters

Parameter	Description
repo_address	Encodes the chosen method and address.
repo_user	Identifies the user name used when logging in to the drivers server. Required for SFTP and FTP methods.
repo_password	The FTP server login password (for user repo_user). Required for FTP method

Table 7 driver_sync.conf parameters (continued)

Parameter	Description
user_private_rsa_key	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.
local_dir	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.
repo_public_rsa_key	The path of the public RSA key of the repository server machine (also known as the host public key).
module_sigcheck	Determines whether to verify (1) or not to verify (0) the downloaded kernel drivers.
emc_public_gpg_key	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 driver_sync.sh is used to synchronize a local repository against a remote one.	
sync_pattern	A regular expression that can be used to fetch several drivers at once. This parameter is not normally needed, and becomes relevant only when driver_sync.sh is used by your own private repository to fetch from a remote (e.g. EMC) repository.
is_local_repo	Set to 0 (default) or omit, for SDC/RF clients. Set to 1 when driver_sync.sh 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 = .*
```

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/XC 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.

Maintaining a ScaleIO system

This section describes how to perform the following maintenance activities to an existing ScaleIO system:

- "Installing RFcache on servers in an existing ScaleIO system"
- "Extending the MDM cluster from 3 to 5-node"
- "Creating a LockBox for SNMP/ESRS"
- "Switching to secured authentication mode"
- "Using SCLI in non-secure mode"
- "Extending an existing ScaleIO system"
- "Removing components"

  Log collection is described in the *EMC ScaleIO Log Retrieval Technical Notes*, which can be downloaded from EMC Online Support.

Installing RFcache on servers in an existing ScaleIO system

This section describes how to install RFcache to servers in an existing ScaleIO system.

Proceed to the section that matches your operating system environment:

- "Installing RFcache in physical servers"
- "Installing RFcache in VMware servers"

Installing RFcache in physical servers

Installing RFcache in physical Linux or Windows servers is performed with the Installation Manager and the CSV topology file, as described in "Extending an existing ScaleIO system".

The only difference between Linux and Windows installation is that 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 *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 VMware 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:
   - If the server was upgraded from ScaleIO v1.32.x.x, copy the following file to all SVMs running an SDS, then install it with the following command:

     ```
     rpm -i EMC-ScaleIO-xcache-x.x-x.0.slesxx.x.*SVM*.x86_64.rpm
     ```

   - If the server was installed as v2.0 or later, copy the following file to all v2.0 SVMs running an SDS, then install it with the following command:

     ```
     rpm -i EMC-ScaleIO-xcache-x.x-x.0.slesxx.x.x86_64
     ```

3. Use the VMware 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.

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"
- "Extending the MDM cluster in VMware servers"

Extending the MDM cluster in physical servers

Before you begin

If you have upgraded your ScaleIO environment from v1.32.x, you must modify the CSV topology file prior to extending the MDM cluster. For more information, see Converting the CSV topology file after upgrade from ScaleIO v1.32.x on page 188.

You can extend the MDM cluster in physical Linux or Windows servers using the Installation Manager and the CSV topology file.

To extend the MDM cluster, perform the following:

Procedure

1. Get one of the following CSV topology files:
   - The converted CSV used to deploy a v1.32.x system that has been upgraded to v2.x.
   - 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 TieBreaker (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 TieBreaker (TB) role.

Note

If you need to change the roles of the current nodes, do so only after extending the cluster.

Figure 5 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 6 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 187.

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 VMware plug-in, described in "Task 2: Updating the SDC parameters".

Note

If an SDC in your system is from ScaleIO pre-v1.32.4, you must contact EMC Support to get an update `drv_cfg` file for your environment.

Extending the MDM cluster

Procedure

1. If the ESX servers to be added do not have the SDC component on them, install the SDC on each of the ESX servers, as described in "Task 3: Installing the SDC on ESX hosts".
2. Begin a new VMware deployment, as described in the *ScaleIO Deployment Guide*.
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 **TieBreaker MDM** fields that now appear, select the nodes to add to the cluster.

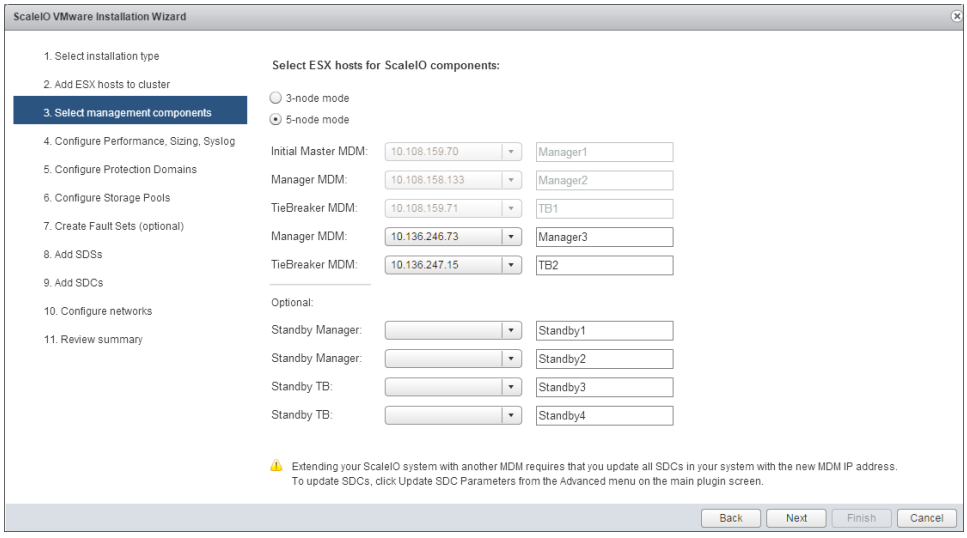

7. Click **Next**, and continue the deployment.

   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
```

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 *EMC 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".

Using SCLI in non-secure mode

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

On every MDM server, disable secure communication:

Procedure

1. Windows

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

2. Linux

   Run this command: `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. Begin a new VMware deployment, as described in the *EMC ScaleIO Deployment Guide*.
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 7 Set Virtual IPs for ScaleIO system screen

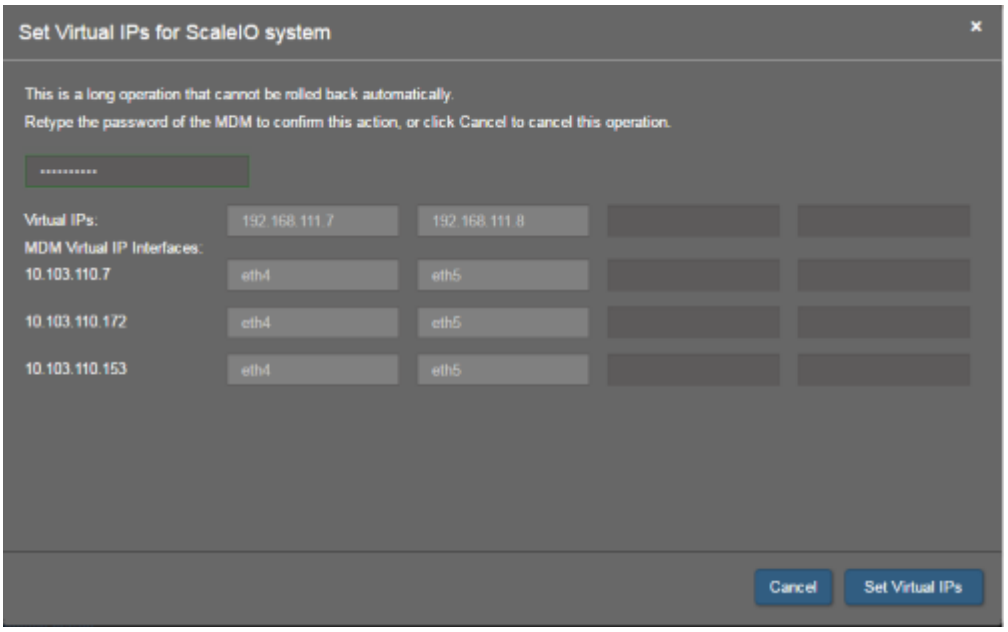

4. For each MDM that you wish to set a virtual IP address, enter a virtual IP address and the NIC to which it will be mapped. For each new virtual IP address, enter the virtual IP address and NIC name for each MDM to which it will be mapped.

   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.

Configuring virtual IP addresses

Configure virtual IP addresses in the vSphere Web plug-in.

Procedure

1. From the **ScaleIO Systems** screen, click **Actions** and select **Configure virtual IPs**.
2. In the **Configure virtual IPs** dialog box, select the network and enter a virtual IP address.

Figure 8 Configure virtual IPs dialog box

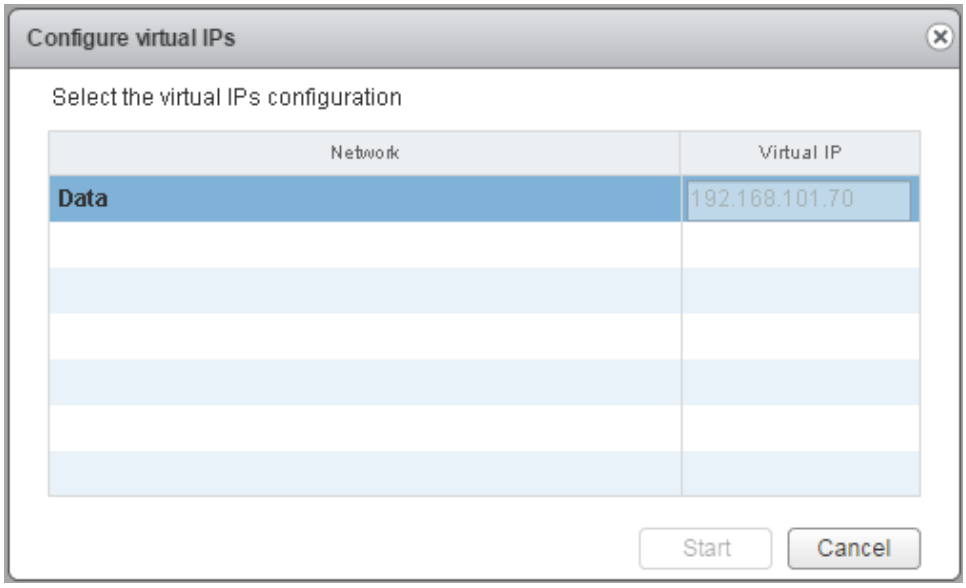

Note

Virtual IP addresses can only be added. To change or remove addresses, use the CLI.

After you finish

Update the SDC parameters to update the SDC configuration. For more information, see Updating SDC parameters on page 197.

Updating SDC parameters

Updating SDC parameters is necessary to ensure MDM-SDC communication when MDM IP addresses have been added or changed. This procedure, performed with the plug-in, is described in the *ScaleIO Deployment Guide*.

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

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**.

Configure ESRS after upgrading

After upgrading, configure EMC Secure Remote Support (ESRS) for remote support.

If ESRS was not enabled for your system prior to the upgrade, following the procedure for configuring ESRS post-installation. For details, see "Configuring ESRS connection properties" in the *EMC ScaleIO Deployment Guide*.

If ESRS was enabled for your system prior to the upgrade, you must perform the following procedures to configure it post upgrade:

1. Enabling ESRS certificate verification on page 199.
2. Adding the ESRS Gateway's certificate to the truststore on page 199.

Before configuring ESRS, ensure that:

- ESRS Gateway v3 version 3.08 or later is installed and configured.
- ESRS Gateway is reachable from ScaleIO on port 9443.
- You have one or more IP addresses of the ESRS servers.
- You have the MDM username and password.

Enabling ESRS certificate verification

Enable certificate verification in the ESRS properties file.

Before you begin

Ensure that a LockBox has already been created and the MDM credentials have been added to it. For more information on creating a LockBox, see the *EMC ScaleIO Deployment Guide*.

Enable the certificate verification feature in the `esrs.properties` file.

Procedure

1. Use a text editor to open the `esrs.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 `DisableSSLVerification` and edit it as follows:

```
DisableSSLVerification=false
```

Results

Certificate verification is enabled.

Adding the ESRS Gateway's certificate to the truststore

Add the ESRS Gateway's certificate to the truststore.

Procedure

1. In a web browser, browse to `<ESRS_Gateway_IP_address:9443>`.
2. Download the certificate that is displayed, and save it as a file.
3. 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
```

4. 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:
   - Windows: Restart the EMC ScaleIO Gateway service.
   - Linux: Run the command `service scaleio-gateway restart`

CHAPTER 12

Troubleshooting

This chapter describes ScaleIO troubleshooting. Topics include:

- Troubleshooting ScaleIO 202
- After SDC installation, the ScaleIO SVM does not start automatically 202
- Application server does not see a ScaleIO volume 202
- Cannot log in to the Installation Manager after upgrade 203
- Certificate error when installing SDC on Windows servers 204
- Enabling acceleration in a new node 204
- Gateway server recovery 207
- Installation Manager returns an error 211
- Installation with the Installation Manager fails 212
- LIA upgrade fails in Windows 2008 212
- Older SDCs cannot communicate with newer SDSs after upgrade to v2.x 212
- Removing RFcache from Windows servers 213
- Removing RFcache leftovers from the Windows OS registry 213
- Replacing a faulty caching device in ScaleIO 213
- Replacing a faulty accelerated storage device in ScaleIO 215
- SVM manual memory allocation 216
- RFcache is not installed on an upgraded Windows server 218
- SDS with RFcache remains in "Disconnected, Decoupled" status after reboot 219
- After retrying upgrade, "Failed to query node" message is displayed 219
- After upgrading, a "Could not switch to cluster mode" error message is displayed 220
- MDM package is not updated during upgrade 220
- Upgrade from v1.32.x causes Data Unavailability status 220
- An SDS performance profile may not get set to High during upgrade 221

Troubleshooting ScaleIO

This chapter describes solutions to issues that may arise. Most of the issues requiring troubleshooting are related to events and CLI return messages:

Event notifications
Upon changes or events in the system, ScaleIO may generate a system event that will be logged in a file. When addressing an event, look for the proper entry in the System Events appendix of the *EMC ScaleIO User Guide*, according to the name of the event received. For each event that requires attention, the entry will include a description of a possible action to take. Follow the instructions before contacting EMC Support.

CLI messages
When issuing a CLI command, ScaleIO generates a textual response describing the outcome of the command. In some cases, the result may be a failure. Upon receiving a failure message, you can address it by looking for the proper entry in the Return Messages appendix of the *EMC ScaleIO User Guide*, according to the text of the message received.

For each return message that requires attention, the entry will include a description of a possible action to take. Follow the instructions before contacting EMC Support.

After SDC installation, the ScaleIO SVM does not start automatically

After deployment is completed, set all ScaleIO VMs to start automatically:

Procedure

1. Click the ESX **Configuration** tab.
2. From the **Software** section, click **Virtual Machine Startup/Shutdown**.
3. Click **Properties**.
4. In the dialog box, select **Allow virtual machines to start and stop automatically with the system**.
5. Select the SVM and move it to the **Automatic Startup** list.
6. Click **OK**.
7. Repeat this process for all SVMs.

Application server does not see a ScaleIO volume

Perform the following steps:

Procedure

1. Check if the ScaleIO system is operational:

```
scli --mdm_ip <mdm IP> --query_all
```

2. Check if the volume is mapped to any of the SDC servers by running the `scli --query_all_volumes` command.

3. Determine if the SDC is installed on the server:
   - Linux: Run `rpm -qa | grep sdc`
   - ESX: Run `esxcli software vib list|grep sdc`
4. Determine if the SDC is connected to an MDM:

```
scli --mdm_ip <mdm IP> --query_all_sdc
```

5. Ensure that the MDM management IP address is up and running.
6. On an application server, rescan for new volumes:
   - Linux: Run `/opt/emc/scaleio/sdc/bin/drv_cfg --rescan`
   - Windows: Run `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg --rescan`
   - ESX: Rescan for new devices

Cannot log in to the Installation Manager after upgrade

If you upgrade the IM, and are then unable to log in, restart the EMC ScaleIO Gateway service.

Certificate error when installing SDC on Windows servers

Due to a Microsoft Windows limitation, when installing an SDC component on a Windows 2008 R2 server, a certificate error, similar to the following may be displayed:

To solve this issue, ensure that Microsoft Security Update KB3033929 is installed.

Enabling acceleration in a new node

Enable acceleration in a new SDS node in a ScaleIO system.

Before you begin

Ensure that you have:

- Admin-level username and password for accessing ScaleIO
- The RFcache package (`xcache`) for installation on a new SDS
- This task requires other ScaleIO procedures, that are described in other documents:
  - Installing an SDS in ScaleIO
  - Adding an SDS to the MDM

- Enabling RFcache, using the Installation Manager (IM)
- Accessing the SVM

Procedure

1. Install an SDS in the ScaleIO system and add it to the MDM.
2. Enable RFcache on the SDS:

SDS OS	Description
ESXi	Access the SVM and then run the following command: `rpm -i EMC-ScaleIO-xcache-2.0-13000.0.sles11.3.*SVM*.x86_64.rpm`
Linux	Enable RFcache using the Installation Manager.
Windows	Enable RFcache using the Installation Manager.

3. Identify the SSD and HDDs to be accelerated and used as cache devices:

SDS OS	Description
ESXi	Run the esxcli commands or open the vSphere View GUI to check the device properties.
Linux	Run the following commands: a. Run the command `fdisk -l` and check the size of the device in the output. b. Run the command `smartctl` and identify the type of device in the output.
Windows	Run either of the following commands: • `wmic diskdrive list brief /format:list` and check the device properties in the output. • `wmic diskdrive get name,size,model` and check the device name, size, and model number in the output.

4. Log in to the Master MDM.
5. Add an SSD device to the default cache pool:

```
scli --add_sds_rfcache_device (--sds_id <ID> |--sds_name
<NAME> | --sds_ip <IP> [--sds_port <PORT>])
--rfcache_device_path <PATH> [--rfcache_device_name <NAME>]
```

Example:

```
scli --add_sds_rfcache_device --sds_name SDS1 --
rfcache_device_path /dev/sdb --rfcache_device_name
rfcache1_SDS1_sdb
```

6. Repeat the previous steps, from the beginning of the procedure, for each new SDS device to be configured in the RFcache Storage Pool.

7. Enable caching on the Storage Pool in the Protection Domain:

```
scli --set_rfcache_usage (((--protection_domain_id <ID> | --
protection_domain_name <NAME>) --storage_pool_name <NAME>) |
--storage_pool_id <ID>) [--use_rfcache | --dont_use_rfcache]
```

Example:

```
scli --set_rfcache_usage --protection_domain_name PD1 --
storage_pool_name SP1 --use_rfcache
```

8. Verify the status of acceleration per SDS:

```
scli --query_sds (--sds_id <ID> | --sds_name <NAME> |--sds_ip
<IP>)
```

Example:

```
scli --query_sds --sds_ip 192.168.1.6
```

Output, similar to the following, appears, for different command types:

- Command to verify the RFcache status:

```
scli --query_sds (--sds_id <ID> | --sds_name <NAME> |--
sds_ip <IP>)
```

Example:

```
scli --query_sds --sds_ip 9.10.21.141
```

Output:

```
SDS 4eecc17300000002 Name: SDS_10.136.215.203 Version:
2.0.13000
Protection Domain: 928cd6e100000000, Name:
Default_Protection_Domain
DRL mode: Volatile
Authentication error: None
IP information (total 2 IPs):
         1: 9.10.21.141      Role: All (SDS and SDC)
         2: 9.10.121.141     Role: All (SDS and SDC)
        Port: 7072
RAM Read Cache information:
        128.0 MB (131072 KB) total size
        Cache is enabled
        RAM Read Cache memory allocation state is PENDING
Rfcache enabled
.
.
.
```

```
Rfcache device information (total 1 devices):
         1: Name: 50000396CC883494  Path: /dev/sdc  Original-
path: /dev/sdc  Size 762495MB  ID: e177c11700020000
```

- Command to verify Storage Pool status:

```
scli --query_storage_pool (((--protection_domain_id <ID> |
--protection_domain_name <NAME>) --storage_pool_name
<NAME>) | --storage_pool_id <ID>)
```

Example:

```
scli --query_storage_pool --storage_pool_name
Cached_Hdd_Pool_1 --protection_domain_name
Default_Protection_Domain
```

Output:

```
Storage Pool Cached_Hdd_Pool_1 (Id: ada1bb8500000000) has 0
volumes and 23.0 TB (23536 GB) available for volume allocation
        The number of parallel rebuild/rebalance jobs: 2
        Rebuild is enabled and using Limit-Concurrent-IO
policy with the following parameters:
        Number of concurrent IOs per device: 1
        Rebalance is enabled and using Favor-Application-IO
policy with the following parameters:
        Number of concurrent IOs per device: 1, Bandwidth
limit per device: 10240 KB per second
        Background device scanner: Disabled
        Zero padding is enabled
        Spare policy: 36% out of total
        Checksum mode: disabled
        Doesn't use RAM Read Cache
        Uses Flash Read Cache
```

9. Repeat this entire procedure on the rest of the SDS nodes.

Results

The new nodes are configured in the RFcache Storage Pool in ScaleIO.

Gateway server recovery

Describes the procedures for recovering a ScaleIO gateway when the physical host server fails:

- Recovering Gateway on a new server during upgrade on page 207
- Recovering IM with LIA trusted IP feature enabled on page 209
- SNMP configuration recovery post gateway crash during upgrade on page 210

Recovering Gateway on a new server during upgrade

Recover the ScaleIO Gateway (GW) on a new server and continue with the process of ScaleIO upgrade, when the GW server fails during ScaleIO upgrade.

If the GW server fails during an upgrade, the Installation Manager (IM) will also go down. You must recover the GW first, and then the IM. This procedure will cause the MDM to use the new GW in place of the failed one.
To recover the Gateway on a new server, perform the following steps:

Procedure

1. Install the Gateway on a new server (identical to the one that crashed), using any of the following options:
   - Linux OS: Use the GW RPM file.
   - Windows OS: Use the GW MSI file.
2. If the LIA-trusted IPs feature was enabled during installation, you must perform the steps described in Recovering IM with LIA trusted IP feature enabled on page 209 before continuing with this procedure.
3. Open the `gatewayUser.properties` file, located at `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/gatewayUser.properties`, and edit the IM parameters.

Parameter	Description	Action
`upgrade.mdm.data.ips`	Configures MDM data IPs	Enter comma-separated IP addresses for different cluster nodes. • If a single node has multiple IP addresses, then enter plus-sign-separated IPs. • Ensure that the sequence of items in the list matches the order of data, management (mgmt), and role for the same node. • Reserve the empty values for Tie Breaker management (TB mgmt) IP addresses. Example: • `upgrade.mdm.data.ips` = 192.100.1.10,192.100.1.49,192.100.1.9,192.100.1.17,192.100.1.11 • `upgrade.mdm.mgmt.ips` = 10.76.60.41,10.76.60.48,10.76.60.41,,
`upgrade.mdm.mgmt.ips`	Configures MDM management IPs	Enter comma-separated IP addresses for different cluster nodes. • If a single node has multiple IP addresses, then enter plus-sign-separated IP addresses • Ensure that the sequence of items in the list matches the order of data, management (mgmt), and role for the same node • Reserve the empty values for Tie Breaker management (TB mgmt) IP addresses. Example: • `upgrade.mdm.data.ips` = 192.100.1.10,192.100.1.49,192.100.1.9,192.100.1.17,192.100.1.11 • `upgrade.mdm.mgmt.ips` = 10.76.60.41,10.76.60.48,10.76.60.41,,
`upgrade.mdm.role`	Configures MDM roles	Enter the role of the MDM server; mdm, slave, or tb (Tie Breaker). The line of the MDM role should correlate to the MDM data IP and MDM management IP addresses. Example: `upgrade.mdm.role` = mdm,slave,slave,tb,tb
`upgrade.mdm.version.target`	Configures the target version of MDM	Enter the MDM target version to which you want to upgrade. The first item in the `version.target` must represent the first cluster node.

Parameter	Description	Action
		Example: `upgrade.mdm.version.target` = 2.0-828.0,2.0-828.0,2.0-828.0
`upgrade.mdm.version.orig`	Configures the original version of MDMs	Enter the MDM's original version, from which you are upgrading. The first item in the `version.orig` must represent the first cluster node. Example: `upgrade.mdm.version.orig` = 1.32-255.0,1.32-255.0,1.32-255.0
`upgrade.mdm.actor.port`	Configures MDM actor ports	Enter comma-separated MDM actor ports, as an ordered list. Each entry is a relevant option, rollback or upgrade. Option selection relates to the operation, version upgrade or rollback that was being performed before the GW crashed. Example: `upgrade.mdm.status` = upgrade
`upgrade.sds.mm_list`	Removes the SDS from Maintenance mode	Change the status of the SDS, which was in maintenance mode before the GW server crash using the `upgrade.sds.mm_list` parameter. Example: `upgrade.sds.mm_list` = 10.76.60.32,10.76.60.33,10.76.60.34,10.76.60.35,10.76.60.36

Results

The IM parameters in the `gatewayUser.properties` file match with the scenario that existed before the GW server crash.

Recovering IM with LIA trusted IP feature enabled

Describes how to recover the Installation Manager (IM) when the LIA-trusted IP feature was enabled in the crashed IM.

When the IM crashes during system upgrade and the LIA trusted IPs feature was enabled with a faulty IM, you are required to update the LIA's white list with the new IM IPs.

Note

When this procedure is required and the Gateway needs to be recovered, you must perform this procedure before recovering the Gateway.

To update the LIA's white list, perform the following steps on all nodes:

Procedure

1. Access the LIA node in the system.
2. Enter the IP addresses:

OS	Steps
Linux	a. Go to: `/opt/emc/scaleio/lia/bin/`

OS	Steps

	b. Run (on one line):
	`./update_conf_bin 7 lia_trusted_ips <Comma_separated_Scaleio_gw_IPs> 1`
	Ensure that you enter all of the ScaleIO GW (IM) hosts' available IP addresses, even if those addresses are not used by ScaleIO.
	Example:
	`./update_conf_bin 7 lia_trusted_ips 1.2.3.4,5.6.7.8 1`

Windows	a. Go to: `C:\Program Files\EMC\scaleio\sds\bin`
	b. Run (on one line):
	`update_conf.exe 7 lia_trusted_ips <Comma_separated_Scaleio_gw_IPs> 1`
	Ensure that you enter all of the ScaleIO GW (IM) hosts' available IP addresses, even if those addresses are not used by ScaleIO.
	Example:
	`update_conf.exe 7 lia_trusted_ips 1.2.3.4,5.6.7.8 1`

SNMP configuration recovery post gateway crash during upgrade

Recover SNMP configuration when the gateway crashes during ScaleIO upgrade.

There are two methods of recovering SNMP configuration:

- Recovering SNMP configuration using Installation Manager on page 210
- Manually recovering SNMP configuration on page 211

Recovering SNMP configuration using Installation Manager

You can use the Installation Manager to recover the SNMP configuration, when the Gateway crashes during a ScaleIO upgrade.

Procedure

1. Delete the alertservice user from the MDM:
   a. Log in to the MDM with the admin user.
      Example:

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

   b. List all existing users:

```
scli --query_users
```

c. Delete the designated user:

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

Example:

```
scli --delete_user --username alertservice
```

d. List all existing users, and ensure that the deleted user is not listed:

```
scli --query_users
```

2. Open the IM GUI and extend the ScaleIO system, using an updated CSV file.
3. In the **Installation Configuration** window, select **Set advanced options (optional)**.

   The **Advanced Configuration** section is displayed.
4. In the **Config. Options** section, select **Enable alert service (required for SNMP and ESRS)**.
5. In the **traps Receiver IP** field, type the traps receiver IP address.
6. Complete the extend operation.

Manually recovering SNMP configuration

Recover SNMP configuration manually, when the gateway crashes during ScaleIO upgrade.

Procedure

1. Create a LockBox.
2. On the Gateway, edit the `gatewayUser.properties` file:

   For more information, see the *EMC ScaleIO User Guide*

   a. Change the value of `features.enable_snmp` to *true*.

   b. Enter IP addresses in `snmp.traps_receiver_ip`.

      ScaleIO supports up to two comma-separated TRAP receivers.

   c. In `snmp.sampling_frequency`, update the MDM sampling period (optional step).

   d. In `snmp.resend_frequency`, update the time period to re-send the already existing TRAPs (optional step).
3. Restart the scaleio-gateway service.

Installation Manager returns an error

In the Installation Manager, if an `HTTP 404 Webpage not available` error (or similar) is displayed, the IM may have been disabled.

For more information, refer to the "Configuring the Installation Manager" section in the *Deployment Guide*.

Installation with the Installation Manager fails

If the following error message appears when installing with the Installation Manager,

```
Command failed: Could not connect to <IP_address>
```

It could be that this node is not in an accessible network.

LIA upgrade fails in Windows 2008

When upgrading or installing the ScaleIO system on Windows 2008R2-based servers, the upgrade or installation may fail and generate this error message: `<package X> for OS windows is not signed`.

Due to an issue in Windows 2008, the certificate issuer might not be recognized. This will result in failure to upgrade or install the LIA MSI in Windows 2008.

To resolve the problem, use the following workaround:

Procedure

1. Download and extract the `roots.zip` root certificates package, located at https://www.symantec.com/theme/roots.
2. Manually install all the certificates in the VeriSign Root Certificates subfolders:
   a. Open a subfolder.
   b. Double-click a certificate, and select **Install Certificate**.
   c. Select **Place all the certificates in the following store**, and then browse to and select the **Trusted Root Certification Authorities** store.
   d. Click **OK**, then **Next**, and then **Finish** .
   e. Review and approve the security warning.

      A message appears, informing you that the import was successful.

   f. Repeat for each certificate.
3. Retry the upgrade or installation.

   The upgrade or installation should now succeed.

Older SDCs cannot communicate with newer SDSs after upgrade to v2.x

After upgrading to v2.x, older SDCs may not be able to communicate with newer SDSs.

If the IP addresses of a v2.x SDS are defined in IPv6 only, older-generation SDCs using IPv4 will not be able to communicate with the SDS. This can cause I/O errors. This can occur when v1.32 SDCs are added to a v2.0 system, as v2.0 supports IPv6, whereas earlier versions do not.

To resolve the problem, if the SDCs are to remain pre-v2.0:

- Retain at least one IPv4 address on the SDS.
- Before removing all the IPv4 IP addresses from the SDSs, ensure that all SDCs are upgraded to v2.x and then rebooted .

Removing RFcache from Windows servers

To remove RFcache from Windows servers, use the Windows Control Panel Add/Remove function to remove xcache.

If SDC is also installed on that server, you must remove SDC first, restart the machine, and then remove xcache.

Removing RFcache leftovers from the Windows OS registry

Run a batch file to remove RFcache leftovers from the Windows registry.

Uninstalling RFcache leaves some leftover entries in the Windows registry. This will cause errors the next time RFcache is installed. To remove the leftover entries in the registry, use the `Clean_XC_registry.bat` file included in the ScaleIO Windows download.

Procedure

1. Locate the `Clean_XC_registry.bat` file that is included in the ScaleIO Windows download.
2. From the Windows command prompt, run the `Clean_XC_registry.bat` file.

Replacing a faulty caching device in ScaleIO

Replace a faulty device that provides acceleration (caching) in a ScaleIO system.

Before you begin

Ensure that you have:

- Administrator authentication credentials to run ScaleIO SCLI commands
- Access to the ScaleIO nodes
- Any of the following parameters:
  - SDS name
  - SDS IP address
  - SDS ID

This procedure applies to ScaleIO software-only systems. It does not apply to AMS-supported ScaleIO systems.

Procedure

1. Log in to the scli interface.
2. Identify the failed RFcache device:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Output, similar to the following, appears for the RFcache devices:

```
Rfcache device information (total 1 devices):
         1: Name: N/A  Path: /dev/sdi  Original-path: /dev/sdi
Size 952719MB  ID: dbfb70fc00000000
```

For a faulty device, the command output displays the fault (error) at the end of the output, while the output for a healthy device does not.

3. Remove the faulty RFcache device from the Storage Pool:

```
scli --remove_sds_rfcache_device (--rfcache_device_id <ID> |
((--sds_id <ID> | --sds_name <NAME> | --sds_ip <IP> [--
sds_port <PORT>]) (--rfcache_device_name <NAME> | --
rfcache_device_path <PATH>)))
```

Example:

```
scli --remove_sds_rfcache_device --rfcache_device_id
7cf1ce1f00000000
```

4. Physically replace the faulty device, using the relevant system and vendor guidelines.
5. Add a new device:

```
scli --add_sds_rfcache_device (--sds_id <ID> |--sds_name
<NAME> | --sds_ip <IP> [--sds_port <PORT>])
--rfcache_device_path <PATH> [--rfcache_device_name <NAME>]
```

Example:

```
scli --add_sds_rfcache_device --sds_name SDS1 --
rfcache_device_path /dev/sdb --rfcache_device_name
rfcache1_SDS1_sdb
```

6. Reset the device error counter:

```
scli --clear_sds_rfcache_error (--sds_id <ID> |--sds_name
<NAME> | --sds_ip <IP> [--sds_port <PORT>])
```

7. Verify the status of acceleration per SDS:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Example:

```
scli --query_sds --sds_name SDS1
```

The command output should display the device state without errors (that is, the replaced device in healthy state); otherwise, repeat the procedure.

Results

The faulty caching device is replaced in the ScaleIO system.

Replacing a faulty accelerated storage device in ScaleIO

Replace a faulty accelerated (cached) storage device in a ScaleIO system.

Before you begin

Ensure that you have:

- Administrator authentication credentials to run ScaleIO SCLI commands
- Access to the ScaleIO nodes
- Any one of the following parameters:
  - SDS name
  - SDS IP address
  - SDS ID

This procedure applies to ScaleIO software-only systems. It does not apply to AMS-supported ScaleIO systems.

Procedure

1. Log in to the scli interface.
2. Identify the failed cached storage HDD device:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Output, similar to the following, appears:

```
Device information (total 7 devices):
         1: Name: N/A  Path: /dev/sdb  Original-path: /dev/sdb
ID: ffcfa55700000000
                Storage Pool: sp1, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         2: Name: N/A  Path: /dev/sdc  Original-path: /dev/sdc
ID: ffcfa55800000001
                Storage Pool: sp1, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         3: Name: N/A  Path: /dev/sdd  Original-path: /dev/sdd
ID: ffcfa55900000002
                Storage Pool: sp1, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         4: Name: N/A  Path: /dev/sde  Original-path: /dev/sde
ID: ffcfa55a00000003
                Storage Pool: sp2, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         5: Name: N/A  Path: /dev/sdf  Original-path: /dev/sdf
ID: ffcfa55b00000004
                Storage Pool: sp2, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         6: Name: N/A  Path: /dev/sdg  Original-path: /dev/sdg
ID: ffcfa55c00000005
                Storage Pool: sp2, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         7: Name: N/A  Path: /dev/sdh  Original-path: /dev/sdh
ID: ffcfa55d00000006
                Storage Pool: sp3, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
```

For a faulty device, the command output displays the fault (error) at the end of the output, while the output for a healthy device does not.

3. Remove the faulty device from the Storage Pool:

```
scli --remove_sds_device (--device_id <ID> |((--sds_id <ID> |
--sds_name <NAME> | --sds_ip <IP> [--sds_port <PORT>]) (--
device_name <NAME> | --device_path <PATH>)))
```

4. Physically replace the faulty device, using the relevant system and vendor guidelines.
5. Add a new device, as a cached device, to the Storage Pool:

```
scli --add_sds_device (--sds_id <ID> | --sds_name <NAME> | --
sds_ip <IP> [--sds_port <PORT>])
--device_path <PATH> [--device_name <NAME>] (--
storage_pool_name <NAME>) | --storage_pool_id <ID>)
```

When the device is added to the Storage Pool, the device starts being cached.

6. Verify the status of the replaced device:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Example:

```
scli --query_sds --sds_name SDS1
```

The command output should display the device state without errors (that is, the replaced device in healthy state); otherwise, repeat the procedure.

Results

The faulty cached device is replaced in the ScaleIO system.

SVM manual memory allocation

When using the plug-in for a clean deployment, SVM memory allocation is performed automatically. In the following cases, SVM memory allocation must be performed manually:

- Manual deployment on VMware.
- Extending an existing SVM with a new ScaleIO role/component, whether this is being done with the plug-in or manually.
  Workaround: Perform all the parts of step 1 and step 2 before extending the additional role/component on the SVM. Perform the steps on one SVM at a time.
- Changing the SDS performance profile, post deployment, or after an upgrade from v1.32.x to v2.0.
  Workaround: Perform all the parts of step 1 one SVM at a time.
- Post backend upgrade from v1.32.x to v2.0.

Workaround: Perform all the parts of step 1 and step 2 on all SVMs that were upgraded, in the following order, one SVM at a time.

- Adding capacity to an SDS that was deployed with partially-populated capacity. Workaround: Perform all the parts of step 1 one SVM at a time.

Procedure

1. For SVMs that are SDS-only, perform the following:
   a. Move the SDS to maintenance mode (MM).
   b. Shut down the SVM.
   c. Increase SVM memory, according to the formula below.
   d. Power up the SVM.
   e. Exit MM.
2. For SVMs that are MDM (Master, Slave, or TB, may contain SDS, also):
   a. Start with Slaves and TBs:
      a. Move the SDS to maintenance mode (MM).
      b. Shut down the SVM.
      c. Increase SVM memory, according to the formula below.
      d. Power up the SVM.
      e. Exit MM.
   b. Proceed with the Master MDM:
      a. Switch ownership, so the Master MDM is now a Slave MDM.
      b. Move the SDS to maintenance mode (MM).
      c. Shut down the SVM.
      d. Increase SVM memory, according to the formula below.
      e. Power up the SVM.
      f. Exit MM.

The memory allocation formula:

Component	Memory allocation rules
Base SVM	• 350 MB
MDM (Master/Slave)	• 470 MB + (500 KB * 8 TB of volume capacity) + (1.44 KB * number of volumes) + (4 KB * number of SDS devices) • Maximum supported volumes: 256 K
Tie Breaker MDM	• 50 MB
SDS	• (Base) 536 MB + (RmCache Size) * 1.15 + (Storage capacity in TB) * 53 MB • For SDS high performance profile, we add 195 MB.

Component	Memory allocation rules		
SDC	• 132 KB + 23 MB * (number of MDMs) + 25 KB * (number of SDSs) + 1.5 KB * (number of volumes) + 16 B * (number of volume blocks) + 24 KB * (8 TB of volume capacity) • Volume blocks: 1 GB storage = 8 volume blocks		
RFcache	• 16 * (cache_size/page_size) • Commonly-used sizes:		
	RFcache page size	RFcache memory requirement, if the cache device is 800 GB	RFcache memory requirement, if the cache device is 1.6 TB
	64 K	200 MB	400 MB
	32 K	400 MB	800 MB
	16K	800 MB	1.6 GB
	8 K	1.6 GB	3.2 GB
	4 K	3.2 GB	6.4 GB

RFcache is not installed on an upgraded Windows server

When using the Installation Manager (IM) to upgrade ScaleIO from v2.x to a later 2.x version on a Windows server, the LIA removes the old RFcache file (`xcache`) and then reboots the host. After the reboot, the LIA fails to install the new RFcache. This error is displayed in the IM.
To avoid this, before starting the upgrade process from the Installation Manager, perform the following steps:

Note

If the SDC is also installed on the Windows server, you must first remove the SDC, restart the machine, and then remove RFcache.

Procedure

1. Log in to the Master MDM, either directly, or via SSH/RDP.
2. On each relevant Storage Pool, disable RFcache:

```
scli --set_rfcache_usage --protection_domain_name <PD_NAME> --
storage_pool_name <SP_NAME> --dont_use_rfcache
```

3. Remove each RFcache device:

```
scli --remove_sds_rfcache_device --sds_ip <SDS_IP> --
rfcache_device_path <RFCACHE_PATH>
```

4. On all Windows hosts with RFcache, perform the following steps, one host at a time:

a. From **Control Panel** > **Add/Remove Programs**, uninstall EMC-scaleIO-xcache.

b. When prompted, reboot the host.

c. When the host has finished rebooting, clean the registry by running the `Clean_XC_registry.bat` script.

The script is included in `Complete_Windows_SW_Download.xip`.

5. Upgrade the Gateway.

   From the extracted download file, copy the ScaleIO Gateway MSI to the IM server and run the file: `EMC-ScaleIO-gateway-2.0-13000.X-x64.msi`.

6. After the Gateway is upgraded, when uploading packages to the Gateway, do not upload the RFcache package.
7. Start the upgrade, as described in the relevant section of the upgrade documentation.
8. Once the upgrade is complete, perform an extend operation to add back RFcache.

SDS with RFcache remains in "Disconnected, Decoupled" status after reboot

When upgrading an SDS server with RFcache (`xcache` package), after rebooting, the RFcache process may load after the SDS process (instead of before the SDS). This results in the SDS remaining in "Disconnected, Decoupled" status after a server reboot.

To fix this issue, restart the SDS process manually, after rebooting.

After retrying upgrade, "Failed to query node" message is displayed

If upgrade via the AMS fails and you restart the upgrade using the GUI **Retry** button, the following error message is returned: `Failed to query node`

Procedure

1. Stop the AMS service.
2. Delete the following file:

   `C:\Program Files\EMC\ScaleIO\AMS\cfg\asyncProcesses.cfg`

3. Start the AMS service.
4. Retry the upgrade.

After upgrading, a "Could not switch to cluster mode" error message is displayed

Upgrades using the AMS GUI might fail, and generate an error message similar to the following: `Command failed: Could not switch to cluster mode on 10.76.59.82, 10.76.59.83 due to: Error`

Procedure

1. Restart the AMS service.
2. In the GUI **Upgrade** tab, click **Close**.
3. Start the upgrade again.

MDM package is not updated during upgrade

After a successful upgrade, the MDM packages may not have been updated to the latest version.

If this happens, run the upgrade again. On the second attempt, the MDM artifacts will be upgraded.

Upgrade from v1.32.x causes Data Unavailability status

When upgrading a ScaleIO v1.32.x system to v2.x, deleting volumes or snapshots while the MDM cluster is being upgraded can cause the new Master MDM to fail to load. This MDM will restart repeatedly, causing a Data Unavailability (DU) status.

Best practice is to avoid deleting volumes or snapshots while the MDM cluster is being upgraded, a part of the upgrade process that is normally very short. If the DU status is encountered, perform the following workaround.

Procedure

1. Delete the shared memory of the new MDM:

MDM OS	Description
Linux	Run this command: `rm -rf /dev/shm/scaleio_mdm_rep_shm`
Windows	a. Run the command prompt as administrator. b. Enter: `net stop shm_server_service` c. Confirm that you want to stop the EMC meta-data manager and EMC data server services.

MDM OS	Description

	d. Start the MDM service:
	`net start mdm_service`
	e. If an SDS is installed on the same machine, start it:
	`net start sds_service`

An SDS performance profile may not get set to High during upgrade

During Installation Manager upgrade, you set SDS performance profile to High, but it is still at default after the upgrade.

This can happen because there is not adequate RAM for the operation.

Solution: Set the performance profile manually, after the upgrade.

CHAPTER 13

Frequently Asked Questions

This chapter describes frequently asked questions regarding ScaleIO. Topics include:

- Install the ScaleIO GUI 224
- Associating ScaleIO volumes with physical disks 224
- Changing default ports 226
- Adding an external SDC to an existing ScaleIO system 227
- Changing the LIA configuration file 229
- Cleaning the ScaleIO VMware environment and performing a clean install 230
- Configuring ScaleIO devices in Linux LVM 232
- Configuring session timeout parameters 232
- Fixing keytool errors 233
- Installing Java on SUSE 12 servers 233
- Mounting ScaleIO 234
- The ScaleIO Gateway web server isn't responding 235
- Upgrading the Gateway when a custom certificate is used 237
- Uploading a new OVA 237
- Using the same data network for different NICs 237
- What to do when the default self-signed certificate expires 237
- Add another IP address subnet to an MDM cluster 238

Install the ScaleIO GUI

You can install the ScaleIO GUI.

Before you begin

- Ensure that the workstation satisfies the requirements described in the "System Requirements" section of the documentation.
- Get the installation file, `ScaleIO GUI for Windows Software Download`, either from the product ISO or the EMC Support Site.

Procedure

1. Install the GUI:

```
EMC-ScaleIO-gui-2.0-13000.X.msi
```

After you finish

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

Associating ScaleIO volumes with physical disks

This section describes how to associate volumes with physical disks.

Contact ScaleIO Customer Support for access to the troubleshooting utility.

To get ScaleIO volume information, run the `scli --query_all_volumes` (or `--query_all` or `--query_volume`) command.

Output similar to the following appears:

```
Query-all-volumes returned 10 volumes
Protection Domain 0728185d00000000 Name: pd1
Storage Pool ad99eaab00000000 Name: default
<No volumes defined>

Storage Pool ad99eaac00000001 Name: sp1
Volume ID: fac22a6300000000 Name: vol0 Size: 152.0 GB (155648 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6400000001 Name: vol1 Size: 400.0 GB (409600 MB) Mapped to 1 SDC Thin-provisioned
Volume ID: fac22a6500000002 Name: vol2 Size: 80.0 GB (81920 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6600000003 Name: vol3 Size: 392.0 GB (401408 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6700000004 Name: vol4 Size: 96.0 GB (98304 MB) Mapped to 1 SDC Thin-provisioned
Volume ID: fac22a6800000005 Name: vol5 Size: 112.0 GB (114688 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6900000006 Name: vol6 Size: 96.0 GB (98304 MB) Mapped to 1 SDC Thin-provisioned
Volume ID: fac22a6a00000007 Name: vol7 Size: 176.0 GB (180224 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6b00000008 Name: vol8 Size: 272.0 GB (278528 MB) Mapped to 1 SDC Thick-provisioned
Volume ID: fac22a6c00000009 Name: vol9 Size: 360.0 GB (368640 MB) Mapped to 1 SDC Thin-provisioned
```

This output shows the Volume ID and name, as well as other volume information.

Volume information - Linux

On the SDC host, run the following command to get the operating system volume information that correlates to the ScaleIO scini device name:

```
ls -l /dev/disk/by-id/ |grep scini
```

Output, similar to the following appears:

```
lrwxrwxrwx 1 root root 12 Aug 25 19:40 emc-vol-62c093a52d14aec7-fac22a6300000000 -> ../../scinia
lrwxrwxrwx 1 root root 12 Aug 25 19:40 emc-vol-62c093a52d14aec7-fac22a6400000001 -> ../../scinic
lrwxrwxrwx 1 root root 12 Aug 25 19:40 emc-vol-62c093a52d14aec7-fac22a6500000002 -> ../../scinib
lrwxrwxrwx 1 root root 12 Aug 25 19:41 emc-vol-62c093a52d14aec7-fac22a6600000003 -> ../../scinie
lrwxrwxrwx 1 root root 12 Aug 25 19:41 emc-vol-62c093a52d14aec7-fac22a6700000004 -> ../../scinid
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6800000005 -> ../../scinif
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6900000006 -> ../../scinig
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6a00000007 -> ../../scinii
lrwxrwxrwx 1 root root 12 Aug 25 19:42 emc-vol-62c093a52d14aec7-fac22a6b00000008 -> ../../scinih
lrwxrwxrwx 1 root root 12 Aug 25 19:43 emc-vol-62c093a52d14aec7-fac22a6c00000009 -> ../../scinij
```

This output shows the scini volume name and the volume ID.

By matching the volume ID in both outputs, you can match the operating system names, sciniX, with the ScaleIO volume name.

For example:

- scinia = fac22a6300000000 = vol0
- scinic = fac22a6400000001 = vol1

Alternatively, run the `sg_inq /dev/sciniX` SCSI query command. The result of this command includes the EMC volume ID at the bottom of the output, as illustrated in the following figure:

```
Vendor identification: EMC
Product identification: ScaleIO
Product revision level: 1.3
Unit serial number: EMC-62c093a52d14aec7-fac22a6300000000
```

Note

The sg3_utils must be installed on the Linux host in order to run this command.

Volume information - Windows

The `sg_inq.exe` file was added to the MSI installation and can be found at `C:\Program Files\EMC\ScaleIO\SDC\diag\`.

Procedure

1. Run the `sg_inq HardiskX` SCSI query command.

   The result of this command includes the EMC volume ID at the bottom of the output.

2. On the MDM, get the ScaleIO volume information:

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

   Output similar to the following is displayed:

   ```
   Retrieved 5 volume(s)
   VOL-ID 6acb988100000000 MDM-ID 0b246c9a755ca3dd
   VOL-ID 6acb988200000001 MDM-ID 0b246c9a755ca3dd
   VOL-ID 6acb988300000002 MDM-ID 0b246c9a755ca3dd
   VOL-ID 6acb988400000003 MDM-ID 0b246c9a755ca3dd
   VOL-ID 6acb988500000004 MDM-ID 0b246c9a755ca3dd
   ```

3. From the Windows command prompt, run this command:

```
wmic diskdrive get deviceid,serialnumber | findstr "EMC"
```

Output similar to the following is displayed:
`\\.\PHYSICALDRIVE13 EMC-0b246c9a755ca3dd-6acb988500000004`

The first part of the output is the disk name. In our example:
`PHYSICALDRIVE13`

The second part is the disk serial number. The last set of the second part (after the dash) is the ScaleIO volume ID. In our example: `6acb988500000004`

After you finish

You can also get the volume ID from the ScaleIO GUI by displaying the **Identity** pane of the volume's properties sheet from **Frontend** > **Volumes**

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, either free these ports or change the ScaleIO ports to another available port.

Table 8 Default ports

Port used by	Port #	File to change	Field to modify (or to add, if it does not exist)	Notes
MDM listener	6611	**Note** Cannot be modified, and must be available		
MDM Cluster member	9011	`/opt/emc/scaleio/mdm/cfg/conf.txt`	`actor_cluster_port=`*<NEW_PORT>*	
SDS listener	7072*x*	`/opt/emc/scaleio/sds/cfg/conf.txt`	`tgt_port=`*<NEW_PORT>*	For cases where multiple SDSs are installed on the same physical server, use ports 7072+*x*, where *x* is the index of the SDS.
LIA listener	9099	`/opt/emc/scaleio/lia/cfg/conf.txt`	`lia_port=`*<NEW_PORT>*	
Gateway (secure)	443	{gateway installation directory} `/conf/catalina.properties`	`http.port=`443	
Gateway (not secure)	80	{gateway installation directory} `/conf/catalina.properties`	`ssl.port=`80	After changing either of the ports, you must restart the service/daemon, as described in The ScaleIO

Table 8 Default ports (continued)

Port used by	Port #	File to change	Field to modify (or to add, if it does not exist)	Notes
				Gateway web server isn't responding on page 235.
SDBG for MDM (Manager)	25620			These ports are used by ScaleIO internal debugging tools to extract live information from the system for debugging purposes.
SDBG for MDM (Tie-Breaker)	25600			
SDBG for SDS	25640			These ports are used by ScaleIO internal debugging tools to extract live information from the system for debugging purposes.
SDBG for Multi-SDS	2564*x*			These ports are used by ScaleIO internal debugging tools to extract live information from the system for debugging purposes. For cases where multiple SDSs are installed on the same physical server, use ports 2564+*x*, where *x* is the index of the SDS.

Adding an external SDC to an existing ScaleIO system

During manual installation, you can install the SDC according to the OS-specific instructions in the following section, and it will be connected to the existing ScaleIO system:

Installing SDC on an ESX server and connecting it to ScaleIO

Install the SDC with the appropriate parameters to connect it to an existing ScaleIO system.

Before you begin

Ensure that you have:

- The virtual IP address or MDM IP address of the existing system
- Login credentials for the SDC
- The appropriate installation packages for the SDC
- Access to the drv_cfg tool. Contact EMC support for access to this tool on ESX.

The following procedure describes installing an external SDC on an ESX server using the esxcli. Alternatively, you can install the external SDC using the vSphere plug-in.

For more information, see "Installing the SDC on ESX hosts" in the *ScaleIO Deployment Guide*.

Procedure

1. On the ESX on which you are installing the SDC, set the acceptance level:

   ```
   esxcli --server=<SERVER_NAME> software acceptance set --
   level=PartnerSupported
   ```

   where *<SERVER_NAME>* is the ESX on which you are installing the SDC.

2. Install the SDC:

   ```
   esxcli software vib update -d "Full Path"
   ```

3. Set the IP address of the MDM:

   ```
   esxcli system module parameters set -m scini -p
   "IoctlIniGuidStr=<XXXXXX> IoctlMdmIPStr=<LIST_VIP_MDM_IPS>"
   ```

   where

   - *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM
   - *<XXXXXX>* is the version

Results

The SDC is installed on the ESX server and is connected to the ScaleIO system.

Installing SDC on a Linux server and connecting it to ScaleIO

Install the SDC with the appropriate parameters to connect it to an existing ScaleIO system.

Before you begin

Ensure that you have:

- The virtual IP address or MDM IP address of the existing system
- Login credentials for the SDC
- The appropriate installation packages for the SDC

The following procedure describes manually installing an external SDC on a Linux server. Alternatively, you can install the external SDC using the Installation Manager. For more information, see "Extending an existing ScaleIO system" in the *ScaleIO Deployment Guide*.

Procedure

1. Install the SDC:
   - RHEL/CentOS /Oracle Linux

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

- Ubuntu

```
MDM_IP=<LIST_VIP_MDM_IPS> dpkg -i <SDC_PATH>.deb
```

- CoreOS

```
MDM_IP=<LIST_VIP_MDM_IPS> ./<LIST_VIP_MDM_IPS>.bsx
```

where

- *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM
- *<SDC_PATH>* is the path where the SDC installation package is located

Results

The SDC is installed on the Linux server and is connected to the ScaleIO system.

Installing SDC on a Windows server and connecting it to ScaleIO

Install the SDC with the appropriate parameters to connect it to an existing ScaleIO system.

Before you begin

Ensure that you have:

- The virtual IP address or MDM IP address of the existing system
- Login credentials for the SDC
- The appropriate installation packages for the SDC

The following procedure describes manually installing an external SDC on a Windows server. Alternatively, you can install the external SDC using the Installation Manager. For more information, see "Extending an existing ScaleIO system" in the *ScaleIO Deployment Guide*.

Procedure

1. On the Windows server on which you are installing the SDC, run:

```
msiexec /i <SDC_PATH>.msi MDM_IP=<LIST_VIP_MDM_IPS>
```

where

- *<SDC_PATH>* is the path where the SDC installation package is located
- *<LIST_VIP_MDM_IPS>* is a comma-separated list of the MDM IP addresses or the virtual IP address of the MDM

Results

The SDC is installed on the Windows server and is connected to the ScaleIO system.

Changing the LIA configuration file

You can change the default behavior of the LIA by editing the configuration file:

- Windows: `C:\Program Files\emc\scaleio\LIA\cfg\conf.txt`
- Linux: `/opt/emc/scaleio/lia/cfg/conf.txt`

  The following are some values relevant to LIA behavior:

```
lia_token=5
lia_enable_install=1
lia_enable_uninstall=1
lia_enable_configure_fetch_logs=1
```

  For example, to restrict which Gateway IP addresses can access the LIA, add those IP addresses to this line in the `conf.txt` file:

```
lia_trusted_ips=<IP_ADDRESS_1>,<IP_ADDRESS_2>
```

  To set this during LIA installation, set the TRUSTED_IPS environment variable. For example:

```
TRUSTED_IPS=1.2.3.4,5.6.7.8 rpm -i lia.rpm
```

Cleaning the ScaleIO VMware environment and performing a clean install

This topic explains how to clean the ScaleIO VMware environment and perform a clean install while using previously defined networks.

Before you begin

Before you begin, unmap and delete any ScaleIO volumes in your system.

If necessary, unregister your ScaleIO system from within the plugin and delete all the ScaleIO SVMs.

Procedure

1. Set to **Run as administrator**, close the existing PowerCLI sessions and open a new one.
2. Using the PS1 script, unregister the plugin.
3. Stop the vSphere web client service:

   VC Linux: `service vsphere-client stop`
4. Delete the contents of the plug-in folder.

   The vSphere web client (Virgo) plug-in folders are located at:

vCenter	Operating system	Path to file
5.x	Windows	C:\ProgramData\VMware\vSphere Web Client\vc-packages\vsphere-client-serenity
	Linux	/var/lib/vmware/vsphere-client/vc-packages/vsphere-client-serenity/

vCenter	Operating system	Path to file
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 devices, perform the following:

Procedure

1. Edit the `/etc/lvm/lvm.conf` file by adding the following line:

```
types = [ "scini", 16 ]
```

2. If only scini devices are to be used, you can add the following filter:

```
filter = [ "a|/dev/scini*|", "r/.*/" ]
```

3. Once configured, the `lvmdiskscan` command should yield results similar to the following:

```
/dev/scinia  [      96.00 GiB] LVM physical volume
/dev/scinib  [     320.00 GiB] LVM physical volume
/dev/scinic1 [      56.00 GiB]
/dev/scinid  [      32.00 GiB]
1 disk
1 partition
2 LVM physical volume whole disks
0 LVM physical volumes
```

4. Continue with normal LVM steps.

Configuring session timeout parameters

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

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

You can modify these parameters, by editing the MDM `conf.txt` file:

- Linux: `/opt/emc/scaleio/mdm/cfg/conf.txt`
- Windows: `C:\Program Files\emc\scaleio\mdm\cfg\conf.txt`

1. To configure maximum session length, edit the value of the `user_session_hard_timeout_secs` parameter. The minimum is 10 seconds, maximum 10 years, and default 8 hours.

2. To configure session idle time, edit the value of the `user_session_timeout_secs` parameter. The minimum is 10 seconds, maximum 3 months, default 10 minutes.
3. After changing the parameters, restart the MDM service (delete and create service) for the changes to take effect.
4. To ensure persistence after MDM restart, make these changes on every MDM.

Fixing keytool errors

Error during rpm installation command

Error message:

```
No keytool path was found. Please pass SIO_GW_KEYTOOL as an argument
to the rpm installation command.
```

If a message similar to this is displayed after executing the rpm command to install the Gateway, add the location of the `/bin/keytool` file on your server to the command.

Example:

```
SIO_GW_KEYTOOL=/usr/lib/jvm/java-1.6.0-openjdk-1.6.0.0.x86_64/jre
rpm -U <gateway_installation_file_name>.rpm
```

Error during rpm upgrade command

Error message:

```
No keytool path was found. Set the environment variable SIO_GW_KEYTOOL
```

If a message similar to this is displayed after executing the rpm command to upgrade the Gateway, add the location of the `/bin/keytool` file on your server to the command.

Example:

```
SIO_GW_KEYTOOL=/usr/java/default/bin/ rpm -U /tmp/EMC-ScaleIO-
gateway-1.32-363.0.x86_64.rpm
```

Installing Java on SUSE 12 servers

Installation of Java is different in SLES-based distributions because SLES uses update-alternatives commands. For SUSE, we use a TGZ file in place of RPM.

To install Java on SUSE 12 servers:

Procedure

1. Untar the TGZ (for example, `jre-8u60-linux-x64.tar.gz`) to `/usr/java`.

   This creates a directory of `/usr/java/jre1.8.0_60/`.

2. Apply the std update-alternatives procedure:

```
/usr/sbin/update-alternatives --install "/usr/bin/java"
"java" "/usr/java/jre1.8.0_60/bin/java" 40
/usr/sbin/update-alternatives --config java
/usr/sbin/update-alternatives --install "/usr/bin/keytool”
"keytool" "/usr/java/jre1.8.0_60/bin/keytool" 40
/usr/sbin/update-alternatives --config keytool
```

Mounting ScaleIO

The exposed ScaleIO volumes are connected to the servers via the network. To configure mounting options of ScaleIO devices, follow the instructions for your operating system.

Use persistent device names, described in full in Associating ScaleIO volumes with physical disks on page 224.

To mount ScaleIO:

Procedure

1. Determine the `/dev/disk/by-id` correlation to `/dev/sciniX`:

```
ls -l /dev/disk/by-id/ |grep scini
```

Output similar to the following appears:

```
lrwxrwxrwx 1 root root 12 Mar 2 05:35 emc-
vol-7ec27ef55b8f2108-85a0f0330000000a -> ../../scinia
lrwxrwxrwx 1 root root 12 Mar 2 05:35 emc-
vol-7ec27ef55b8f2108-85a0f03200000009 -> ../../scinib
lrwxrwxrwx 1 root root 12 Mar 2 05:35 emc-
vol-7ec27ef55b8f2108-85a0f02c00000003 -> ../../scinic
```

2. Run the mount command:

```
mount /dev/disk/by-id/<EMC-vol-id>
```

Example:

```
mount /dev/disk/by-id/emc-
vol-7ec27ef55b8f2108-85a0f0330000000a /mnt_scinia
```

3. To make the mount command persistent, edit the `/etc/fstab` file according to the instructions for your operating system:
   - RHEL 6.x:
     a. In `/etc/fstab`, use a text editor to add the ScaleIO mount lines:

```
/dev/disk/by-id/emc-
vol-7ec27ef55b8f2108-85a0f0330000000a /mnt_scinia ext4
defaults 0 0
```

b. In `/etc/rc.local`, use a text editor to add the mount commands:

```
mount /mnt_scinia
```

- RHEL 7.x:
  In `/etc/fstab`, use a text editor to add `_netdev` to the ScaleIO mount lines.

  Example:

  ```
  /dev/disk/by-id/emc-vol-7ec27ef55b8f2108-85a0f0330000000a /
  mnt_scinia ext4 defaults,_netdev 0 0
  ```

  Ensure that you comply with the `netdev` and syntax rules for your file system, as described in the `man` page.
- SLES:
  In `/etc/fstab`, use a text editor to add `nofail` to the ScaleIO Ready Node mount lines.

  Example:

  ```
  /dev/disk/by-id/emc-vol-7ec27ef55b8f2108-85a0f0330000000a /
  mnt_scinia ext3 nofail 0 0
  ```

  Ensure that you comply with the `nofail` and syntax rules for your file system, as described in the `man` page.

The ScaleIO Gateway web server isn't responding

The ScaleIO Gateway (REST service, Installation Manager) may be disabled:

The ScaleIO Gateway seems to be locked or disabled, and returns the HTTP status code 401 or 403.

Solution

- Ensure that the Gateway is enabled, as described in the documentation.
- In the `gatewayUser.properties` file, ensure that the `gateway-admin.password` property has a non-blank password. If the password is blank, the gateway has been locked.

  The following table shows the location of the gatewayUser.properties file:

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

To reset the Scaleio-Gateway password, perform the following steps:

Procedure

1. Use SioGWTool to reset the password by typing the following command:

```
SioGWTool --reset_password --password <new_scaleio-
gateway_password> --config_file
<path_to_file_gatewayUser.properties>
```

Note

The path to SioGWTool is:

Linux: `/opt/emc/scaleio/gateway/bin/SioGWTool.sh`

Windows: `C:\Program Files\EMC\ScaleIO\Gateway\bin\SioGWTool.bat`

---

2. Restart the scaleio-gateway service

The ScaleIO Gateway web server isn't responsive and the following error appears in the catalina log file:

- Windows:

  `C:\Program Files\EMC\ScaleIO\Gateway\logs\catalina.<date>.log`

- Linux:

  `/opt/emc/scaleio/gateway/logs`

```
2014-06-21 22:50:57,113 [main] ERROR
o.a.coyote.http11.Http11NioProtocol - Failed to initialize end
point associated with ProtocolHandler ["http-nio-443"]
java.net.BindException: Address already in use: bind
```

Solution

Perform one of the following:

Procedure

1. Find the service/daemon that is currently occupying that port and stop it:
   - Windows

     Run: netstat -anb
   - Linux

     Run: `netstat -alp`

   On Windows, one of the common applications that occupies this port is the VMware workstation, which uses this port for the shared VM feature. You can configure VMware workstation to use a different port via the Settings dialog, or you can disable the shared VM feature.

   Once the port is free, restart the scaleio-gateway service:

   - Windows

     Restart the EMC ScaleIO Gateway service.

- Linux

  Type the command `service scaleio-gateway restart`

2. Change the ScaleIO Gateway web server to run on a different port, as described in "Changing default ports".

   After doing so, restart the ScaleIO Gateway service/daemon, as described above. Access the Gateway with the new port. For example: https://<host>:<port>

Upgrading the Gateway when a custom certificate is used

If a custom security certificate is used on the ScaleIO Gateway (Windows and Linux environments), you must save a copy of the certificate (`*.keystore` file) and the `catalina.properties` file before you upgrade the gateway. After the upgrade is complete, you must copy these files back to their original location.

The default file locations, per operating system, are:

Linux:

`/opt/emc/scaleio/gateway/conf/catalina.properties`

`/opt/emc/scaleio/gateway/conf/certificates/.keystore`

Windows (64 bit):

`C:\Program Files\EMC\ScaleIO\Gateway\conf\catalina.properties`

`C:\Program Files\EMC\ScaleIO\Gateway\conf\certificates\.keystore`

Uploading a new OVA

If you have already used the OVA to create a template, you cannot create another template with the same name in the same datacenter.

Either remove the original template first, or use the `ScaleIOPluginSetup-2.0-13000.X.ps1` script, option #3, to assign a different name to the new template.

You can also upload the OVA manually using the VMware OVA upload tools. Configure the networks manually, after deployment or during the wizard menus. For more information, see the VMware user guides.

Using the same data network for different NICs

This configuration is supported, but it could reduce efficiency of outgoing communication and deny you the benefits of high availability of the multiple networks.

What to do when the default self-signed certificate expires

If the default self-signed security certificate is used on the ScaleIO Gateway, it expires after approximately one year. When you upgrade the gateway, the self-signed certificate is automatically replaced with a new one. If your self-signed security certificate expires, you can create a new one using the Java keytool utility.

Add another IP address subnet to an MDM cluster

Add an IP network to an existing MDM cluster.

Before you begin

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

Note

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

Procedure

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

   ```
   scli --query_cluster
   ```

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

2. Switch to single cluster mode:

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

3. Remove the standby MDM:

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

4. Remove the Tie Breaker:

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

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

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

   For example:

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

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

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

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

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

For example:

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

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

```
scli --query_cluster
```

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

9. Switch MDM ownership to verify cluster functionality:

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

For example:

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

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

```
scli --query_cluster
```

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

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

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

CHAPTER 14

DTK - Hardware Update Bootable ISO

The Dell OpenManage Deployment Toolkit (DTK) includes a set of utilities, sample scripts, and sample configuration files that you can use to deploy and configure the Dell systems. This section provides additional information on using the ScaleIO Ready Node Hardware Update Bootable ISO ("Hardware ISO") to update drivers, BIOS, and firmware on ScaleIO Ready Node servers.

- Dell OpenManage DRAC Tools (RACADM) 242
- SATADOM firmware update 244
- Recommended BIOS and firmware settings 245
- Troubleshooting the Hardware ISO 248
- Additional resources 250
- Setting up the BMC (iDRAC) IP address and BIOS 251
- Verifying status of system hardware, storage controller, and disks 253

Dell OpenManage DRAC Tools (RACADM)

An integrated Dell Remote Access Controller (iDRAC) with Lifecycle Controller is embedded in every Dell PowerEdge server. The RACADM command-line utility provides a scriptable interface that enables you to configure the iDRAC either locally or remotely. The utility runs on the management station and the managed system.

The RACADM utility supports the following interfaces:

- Local - Supports running RACADM commands from the managed server's operating system. To run local RACADM commands, install the OpenManage software on the managed server. Only one instance of Local RACADM can be executed on a system at a time.
- SSH or Telnet (also known as Firmware RACADM) - Firmware RACADM is accessible by logging in to iDRAC using SSH or Telnet.
- Remote - Supports running RACADM commands from a remote management station such as a laptop or desktop running Windows or Linux. To run Remote RACADM commands, install the DRAC Tools utility from the OpenManage software on the remote computer.

Note

For information regarding RACADM commands, see the RACADM CLI Guide.

Updating the hardware using remote RACADM

You can install and execute the Dell RACADM tool from any management system with access to the iDRAC network. The remote RACADM command set is useful in this situation to mount and execute the Hardware ISO to a large number of ScaleIO Ready Node servers. This prevents the need for multiple web browser windows and manual keyboard and mouse clicks for each server.

Before you begin

Note

For additional instructions regarding version updates following SATADOM replacement, see SATADOM firmware update on page 244.

Note

Use calvin as the password for Dell-supplied nodes. For EMC-supplied nodes, the password is password.

Procedure

1. Configure a laptop or server with access to the iDRAC network.

   For details, see Setting up the BMC (iDRAC) IP address and BIOS on page 251.

2. Install the Dell DRAC tools, including RACADM:

   a. In any Internet browser, go to the RACADM Command Line Interface for DRAC page in the Dell Systems Management wiki.

   b. Follow the instructions for downloading and installing Remote RACADM for your operating system.

3. Download the ScaleIO Ready Node Hardware Update Bootable ISO from the ScaleIO Ready Node page and make it accessible on the network share folder.
4. Create a CIFS or NFS network share that includes the Hardware ISO. The share will need to be accessible to the iDRAC network.

   Refer to OS vendor guidelines for instructions on how to create CIFS (windows) or NFS (Linux) shares
5. Open a terminal/command prompt with superuser/administrator accessibility. This prompt should allow execution of the RACADM utility.
6. Mount the Hardware ISO to the iDRAC from the remote share:

```
racadm -r <dracIP> -u root -p calvin remoteimage -c -u
<myuser> -p <mypass> -l //<myip>/SIORN/ScaleIO-Ready-Node-
Hardware-Update-for-Dell_2.0.1.iso
```

   Where:

   - *<dracIP>* is the iDRAC IP address
   - *<myuser>* is the NFS/CIFS share username
   - *<mypass>* is the NFS/CIFS share password
   - *<myip>* is the NFS/CIFS server IP

7. Enable the iDRAC VirtualCD to boot once using these Remote RACADM commands:

```
racadm -r <dracIP> -u root -p calvin set
iDRAC.ServerBoot.BootOnce Enabled
racadm -r <dracIP> -u root -p calvin set
iDRAC.ServerBoot.FirstBootDevice VCD-DVD
```

8. Verify that the server is in a state in which it can be power-cycled. (The OS should be in maintenance mode, as needed).
9. Power-cycle the ScaleIO Ready Node server:

```
racadm -r <dracIP> -u root -p calvin serveraction
graceshutdown
```

10. Repeat steps 6 on page 243 - 9 on page 243 for each ScaleIO Ready Node server needing the BIOS and firmware updates and configuration. All servers can be updated in parallel.
11. Wait for the configuration and firmware updates to complete. The server console screen will indicate when the script is complete.

   WARNING

   Do not reboot the ScaleIO Ready Node server while the update process is being performed!

   The iDRAC will be reset several times during the update process. This causes the iDRAC virtual console viewer to close, virtual media to disconnect, and the iDRAC browser window to be unavailable for several minutes during each reset. The hardware update scripts will continue to run from RAM on the server.

The update script will generate a log indicating whether each configuration and firmware flash is successful.

12. (Optional) Check each ScaleIO Ready Node server's log for successful competition:
    a. Connect to the iDRAC KVM console screen.

       For instructions, see Opening the KVM console on page 127.

    b. After the update script completes, press Alt+F2 to access a user console, and then Enter to log in.
    c. Check the log contents for errors:

       ```
       less /bundleapplicationlogs/apply_components.log
       ```

       For more information, see Troubleshooting the Hardware ISO on page 248 .

    d. Press q to exit the log viewer, and then Alt+F1 to access the original console screen.
13. Reboot the servers and allow the update and configuration jobs to complete. You can do this remotely using RACADM (see Step 9 on page 243) or locally from the server console screen.
14. For each server, check the iDRAC job queue to verify that the iDRAC jobs created by the bootable ISO have completed successfully:

    ```
    racadm -r <dracIP> -u root -p calvin jobqueue view
    ```

    Any job failures may require re-running the bootable ISO, or further troubleshooting.

15. For each server, after the hardware updates are finalized, clean up the iDRAC job queue:

    ```
    racadm -r <dracIP> -u root -p calvin jobqueue delete --all
    ```

16. Continue with OS installation and configuration, ScaleIO deployment, and other required tasks.

    At the end of the process, the BIOS password changes to emcbios.

SATADOM firmware update

The Dell factory should provide the SATADOM boot device with the correct firmware already loaded. However, in some cases, it may be necessary to update the firmware loaded on the SATADOM boot device.

For convenience, SATADOM firmware version S130710K is included on the ScaleIO Ready Node Hardware ISO.

WARNING

The SATADOM firmware update process permanently erases all data from the SATADOM device. Proceed with caution!

Procedure

1. Connect to the iDRAC Virtual KVM console.

   For instructions, see Opening the KVM console on page 127.

2. Boot the ScaleIO Ready Node server to the Hardware ISO.

   For instructions, see Upgrading the firmware using the iDRAC virtual console on page 130.

3. After the update script completes, press Alt-F2 to access a user console.
4. Press Enter to login.
5. Execute the SATADOM firmware:

```
/tmp/S130710K_read_log_issue/satadom.sh
```

6. Press y to begin the SATADOM update.
7. Press Alt+F1 to access the original console screen, and then press Enter to reboot the server.

Recommended BIOS and firmware settings

This section describes the BIOS, firmware, and configuration settings included in the Hardware ISO. Topics include:

- ScaleIO ID module on page 245
- BIOS and firmware on page 245
- Configuration settings on page 246

ScaleIO ID module

The Hardware ISO runs a script that automatically flashes the ID module, as needed.

This action rebrands a PowerEdge server naming, and is intended to make it compatible with the ScaleIO AMS software deployment.

The ScaleIO ID module should be installed only on ScaleIO Ready Node servers. In the event of a system board failure, the Hardware ISO can assist with the reinstallation of the ID module on the replacement system board.

BIOS and firmware

The Hardware ISO runs a script that automatically forces the server to install the firmware listed in the table below.

This firmware is consistent with the qualified ScaleIO Ready Node Driver and Firmware Matrix.

Table 9 Hardware ISO contents

Description	Version
Drivers for OS Deployment Application	16.03.00 A00
Expander Backplane Firmware	3.03 A00
PowerEdge Server BIOS	2.2.5
Backplane Firmware	2.23 A00
SAS HBA330 Storage Controller Firmware	13.17.03.00 A01
PERC H330 Storage Controller Firmware	25.4.0.0017 A04
PERC H730 Storage Controller Firmware	25.4.0.0017 A06
Intel x520/i350 NIC Firmware	17.0.12 A00
iDRAC with Lifecycle Controller Firmware	2.40.40.40 A00

Note

Some of the firmware listed in the table is dependent on the ScaleIO Ready Node hardware configuration. The Hardware ISO attempts to apply all firmware updates, but only those updates that are compatible will be installed.

Applying settings using RACADM

The individual firmware files are also available on the Dell Support website, and can easily be installed using the following remote RACADM command:

```
racadm -r <dracIP> -u root -p calvin update -f <filename.exe>
```

Where:

- *<dracIP>* is the iDRAC IP address
- *<filename.exe>* is the name of the Dell Windows update packages

Note

Use calvin as the password for Dell-supplied nodes. For EMC-supplied nodes, the password is password.

Configuration settings

The Hardware ISO runs a script that automatically configures the BIOS and iDRAC settings listed in the table below. Some settings are dependent on the server model.

Table 10 Hardware ISO configuration settings

Description	Setting	Value
Server BIOS Boot Sequence	BIOS.BiosBootSettings.BootSeq	HardDisk.List.1-1
Hard Disk Boot Order	BIOS.BiosBootSettings.HddSeq	Disk.SATAEmbedded.J-1,RAID.Integrated.1-1

Table 10 Hardware ISO configuration settings (continued)

Description	Setting	Value
		or Disk.SATAEmbedded.D-1,RAID.Integrated.1-1
Server Boot Mode	BIOS.BiosBootSettings.BootMode	Bios
SRIOV Global Enablement	BIOS.IntegratedDevices.SriovGlobalEnable	Enabled
Memory Performance Tuning	BIOS.MemSettings.SnoopMode	EarlySnoop
System Change Tracking	BIOS.MiscSettings.InSystemCharacterization	Disabled
CPU Virtualization Features	BIOS.ProcSettings.ProcVirtualization	Enabled
CPU Cores	BIOS.ProcSettings.ProcCores	All
CPU X2APIC Mode	BIOS.ProcSettings.ProcX2Apic	Enabled
CPU Turbo Engagement	BIOS.ProcSettings.ControlledTurbo	Disabled
System Power Profile	BIOS.SysProfileSettings.SysProfile	PerfOptimized
OS to iDRAC Pass-through Mode	iDRAC.OS-BMC.PTMode	usb-p2p
OS to iDRAC Enablement	iDRAC.OS-BMC.AdminState	Enabled
iDRAC DHCP Enablement	iDRAC.IPv4.DHCPEnable	Disabled
iDRAC Default Credentials	iDRAC.Tuning.DefaultCredentialWarning	Disabled
iDRAC IPMI Enablement	iDRAC.IPMILan.Enable	Enabled
iDRAC Alert Enablement	iDRAC.IPMILan.AlertEnable	Enabled
iDRAC IPv6 Enablement	iDRAC.IPv6.Enable	Enabled

Applying settings using RACADM

The individual settings can also be applied using the remote RACADM command:

```
racadm -r <dracIP> -u root -p calvin set <setting> <value>
```

Where:

- *<dracIP>* is the iDRAC IP address
- *<setting>* is the BIOS/iDRAC setting name
- *<value>* is the BIOS/iDRAC setting value

Note

When setting the BIOS configuration, include this command:

```
racadm -r <dracIP> -u root -p calvin jobqueue create BIOS.Setup.1-1 <value>
```

Troubleshooting the Hardware ISO

This section describes troubleshooting procedures for problems you may encounter while using the Hardware ISO.

Topics include:

- General iDRAC failures on page 248
- iDRAC virtual console issues on page 248
- iDRAC virtual media issues on page 249
- Bootable ISO logs on page 249

General iDRAC failures

When problems occur with iDRAC or Lifecycle Controller jobs, you can delete all jobs with a single iDRAC command. All of the completed jobs, plus any orphaned pending jobs, are deleted, and the data manager service on the iDRAC is restarted.

Procedure

1. Clear the iDRAC job queue:

   ```
   racadm -r <dracIP> -u root -p calvin jobqueue delete -i
   JID_CLEARALL_FORCE
   ```

   Note

   Use calvin as the password for Dell-supplied nodes. For EMC-supplied nodes, the password is password.

2. Wait 120 seconds.

   The iDRAC is unable to process any other jobs during this time.

3. Reset the iDRAC

   ```
   racadm -r <dracIP> -u root -p calvin racreset
   ```

   The iDRAC becomes accessible on the network 3 to 5 minutes after the reset.

iDRAC virtual console issues

Within the iDRAC virtual console window, if the keyboard or some of the keys are not responding, perform the following checks:

- Ensure on your keyboard that the Scroll Lock button is off.

- In the iDRAC virtual console, ensure that the **Keyboard/Mouse Attach State** is set to**Auto-attached**.

For persistent iDRAC virtual console issues, see General iDRAC failures on page 248.

iDRAC virtual media issues

The following are solutions to problems that may occur when using iDRAC virtual media:

1. If the iDRAC is stuck, one of these actions may assist with recovery:
   - In the server's **Attached Media** screen, disconnect the iDRAC Remote File Share from the iDRAC browser.
   - Disconnect the iDRAC remote image:

     ```
     racadm -r <dracIP> -u root -p calvin remoteimage -d
     ```

     Where:

     - *<dracIP>* is the iDRAC IP address

   Note

   Use calvin as the password for Dell-supplied nodes. For EMC-supplied nodes, the password is password.

2. Issues when booting to iDRAC virtual media may cause CPU machine check errors at POST. When this occurs, clear the iDRAC job queue and reset the iDRAC, as described in General iDRAC failures on page 248.

Bootable ISO logs

You can view ScaleIO Ready Node Hardware Update Bootable ISO logs after the update script completes.

Procedure

1. Press Alt-F2 to access a user console, and then press Enter
2. Open the log to check the contents for errors:

   ```
   less /bundleapplicationlogs/apply_components.log
   ```

3. You can also view the script for the Hardware ISO, which is useful in helping to identify and troubleshoot log entries:

   ```
   less /opt/dell/toolkit/systems/drm_files/apply_bundles.sh
   ```

Results

The script attempts to configure several boot order commands, regardless of the hardware configuration of the ScaleIO Ready Node server. This allows the script to support multiple hardware platform configurations.
Therefore, it is normal to see these error messages within the logs:

```
Apply R730xd bootorder
[Key=BIOS.Setup.1-1#BiosBootSettings]
RAC1017: Successfully modified the object value and the change is in
       pending state.
       To apply modified value, create a configuration job and reboot
       the system. To create the commit and reboot jobs, use "jobqueue"
       command. For more information about the "jobqueue" command, see RACADM
       help.

ERROR: BOOT016: Input source argument value for the boot device is incorrect or
        not found among the boot devices on the system.

ERROR: BOOT016: Input source argument value for the boot device is incorrect or
        not found among the boot devices on the system.

ERROR: BOOT016: Input source argument value for the boot device is incorrect or
        not found among the boot devices on the system.

[Key=BIOS.Setup.1-1#BiosBootSettings]
RAC1017: Successfully modified the object value and the change is in
       pending state.
:
```

Firmware updates may also display and log the following message:

```
This update is not compatible with your system configuration.
```

```
iDRAC Settings Complete
Proceeding with server BIOS and firmware flash.
Collecting inventory...
Running validation...

This Update Package is not compatible with your system configuration.

Collecting inventory...
Running validation...
```

These firmware are included on the ISO in order to support various ScaleIO Ready Node platforms. The log messages do not necessarily indicate a failure.

Additional resources

This section contains information regarding additional resources that may be helpful useful when using the Hardware ISO. Topics include:

- ScaleIO resources on page 250
- Dell Lifecycle Controller (LC) on page 251
- Dell OpenManage Deployment Toolkit (DTK) on page 251

ScaleIO resources

ScaleIO Ready Node deployments have specific guidelines regarding server installation, rack and stack procedures, and power and networking requirements.

The *ScaleIO Ready Node Hardware Installation Guide* describes how to install the physical components of a ScaleIO Ready Node system. For additional information regarding the ScaleIO Ready Node product, documentation, advisories, downloads, and white papers, visit the ScaleIO Ready Node product page.

Dell Lifecycle Controller (LC)

With the launch of the Dell PowerEdge 13th-generation servers in September 2014, Dell has enhanced our embedded management without the need to install a software-based agent within the host operating system.

At the heart of the 13th-generation servers' embedded management is the iDRAC8 with Lifecycle Controller (LC) technology. This technology allows users to perform useful tasks such as configuring BIOS and hardware settings, deploying operating systems, updating drivers, changing RAID settings, and saving hardware profiles. Together, they provide a robust set of management functions that can be leveraged throughout the entire server lifecycle.

For more information, visit the Lifecycle Controller wiki homepage.

Dell OpenManage Deployment Toolkit (DTK)

The Dell OpenManage Deployment Toolkit (DTK) includes a set of utilities, sample scripts, and sample configuration files that you can use to deploy and configure Dell systems.

You can use the DTK to build script-based and RPM-based installation for deploying large number of systems on a pre-operating system environment in a reliable way, without changing their current deployment processes. Using DTK you can install operating systems on Dell systems in BIOS or Unified Extensible Firmware Interface (UEFI) mode.

For more information, visit the Dell OpenManagement Deployment Toolkit wiki homepage.

Setting up the BMC (iDRAC) IP address and BIOS

Set up the BMC (iDRAC) IP address and set up or validate the BIOS on the ScaleIO Ready Node servers.

Before you begin

Ensure that you have access to, or have the details for:

- The KVM console
- The server BMC (iDRAC) IP address
- The server BMC (iDRAC) subnet mask
- The Gateway IP address
- The VLAN ID of the BMC (iDRAC), if the VLAN is used

Note

The BMC (iDRAC) IP address may be an IPv4 or an IPv6 address.

Defaults:

- BMC (iDRAC) username - root
- BMC (iDRAC) password - password
- BIOS password - emcbios

During the BMC (iDRAC) IP and BIOS setup, use the following operations:

- Use the arrow keys to navigate in the BIOS screens.
- Use the + and - keys on the keyboard to change the option selection in the BIOS screens.
- Use SPACE or ENTER keys to change the settings.
- The ENTER key opens a list to the desired values.

Procedure

1. Open the KVM console.

   For console operations (KVM access), ensure that you have either a VGA tool kit/Crash Cart to allow physical console connection from a laptop computer to a server, or a computer screen and keyboard connection to the rack.
2. Press F2 to access the main menu.
3. From the **System Setup Main Menu** screen, select the **iDRAC Settings** menu option.
4. Configure network settings:
   a. In the **iDRAC Settings** screen, select **Network**.
   b. In the **iDRAC Settings--Network** pane, verify the following parameter values:
      - **Enable NIC** = **Enabled**
      - **NIC Selection** = **Dedicated**
   c. From the **IPv4 Settings** pane, configure the IPv4 parameter values for the BMC (iDRAC) port:
      - **Enable IP IPv4** = **Enabled**
      - **Enable DHCP** = **Disabled**
      - **Static IP Address** = Static IP address
      - **Static Gateway** = Gateway IP address
      - **Static Subnet Mask** = Subnet mask IP address
   d. From the **IPv6 Settings** pane, configure the IPv6 parameter values for the BMC (iDRAC) port.
   e. From the **IPMI Settings** pane, verify the following parameter values:
      - **Enable IPMI Over LAN** = **Enabled**
      - **Channel Privilege Level Limit** = **Administrator**
   f. If you are working in a VLAN setup, access the **VLAN Configuration** pane and configure the VLAN ID parameters.
   g. When the parameter set up is complete, click **Back** to display the **iDRAC Settings** screen.
5. From the **iDRAC Settings** screen, click **Finish**, **Yes**, and then **OK** to return to the **System Setup Main Menu** screen.
6. In the **System Setup Main Menu** screen, select the **System BIOS** menu.
7. In the **System BIOS Settings** screen, verify that the processor settings are correct.

   If the settings are incorrect, configure them as follows:

a. Select **Processor Settings**.

b. In the **System BIOS Settings--Processor Settings** pane, verify the following parameter values:

- **Virtualization Technology** = **Enabled**
- **Number of Cores Per Processor** = **All**

c. Click **Back** to return to the **System BIOS Settings** screen.

8. Configure boot settings,:

a. Select **Boot Settings**.

b. In the **System Boot Settings--Boot Settings** pane, verify that **Boot Mode** is set to **BIOS**.

c. Select the **BIOS Boot Settings** link.

d. In the **System BIOS Settings--Boot Settings--BIOS Boot Settings** pane, verify that:

- In the **Boot Sequence** list, **Hard drive C:** appears as the first item.
- In the **Hard-Disk Drive Sequence** list, the **SATADOM-ML 3SE** device appears as the first item.

e. Click **Back** twice to return to the **System BIOS Settings** screen.

9. Configure integrated devices:

a. Select **Integrated Devices**.

The **System BIOS Settings--Integrated Devices** screen appears.

b. Set **SR-IOV Global Enable** to **Enabled**.

c. Verify that the **Internal USB Port** parameter is set to **Off**.

d. Click **Back** to return to the **System BIOS Settings** screen.

10. From the **System BIOS Settings** screen, click **Finish**, **Yes**, and then **OK** to return to the **System Setup Main Menu** screen.

11. Select **Finish** to exit the BIOS and apply all settings post boot.

Results

The BMC (iDRAC) IP and server BIOS address configuration is complete.

Verifying status of system hardware, storage controller, and disks

Verify the status of the system hardware, storage controller, and disks in a ScaleIO Ready Node server.

Before you begin

Ensure that you know:

- The IP address of the BMC (iDRAC) port
- The username and password for the BMC (iDRAC) portal (default username and password are root and password)

Procedure

1. From a browser, go to http://*<BMC/iDRAC_IP_address>*.

   The **DELL Console Login** window is displayed.

2. Type the user name and password, then click **Login**.

   The **System Summary** dashboard displays high-level status of all HW devices in the **Server Health** pane.
   In an ideal scenario, all hardware sensors should appear in green.

3. To view the event log (SEL log), select **View Logs** in the **Quick Launch Tasks** pane.

   The **System Event Log** pane is displayed with color-coded severity levels.

4. Ensure that any power supply- and fan-related events in the event log are non-repetitive. Repetitive events may be due to the intermittent nature of faults, such as poor physical connections.

   For repetitive events, it is recommended that you remove the relevant hardware module and replace it in its socket.

5. In the navigation pane, select **Storage** > **Controllers**.

6. In the **Health and Properties** pane, a table displays the controller-related information, with controller type in the **Name** column, as follows:

   ```
   PERC H730 Mini (Embedded)
   ```

7. In the navigation pane, select **Storage** > **Physical Disks**.

8. In the **Health and Properties** pane, a table displays the information on physical disks.

9. Verify that no disk is in the Failed state.

   If any of the disks is failed, refer to the relevant disk FRU.

10. In the navigation pane, select **Server**.

11. In the **System Inventory** pane, verify that the server drivers and firmware in the **Firmware Inventory** list match the required versions, as published in the ScaleIO Ready Node Driver and Firmware Matrix.

    If the driver and firmware versions do not match the matrix, you must update them. See Updating the BIOS, firmware and settings on page 128.

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 Random access electronic storage used to retain frequently used data for faster access by the channel.

ScaleIO can be configured to use server RAM cache, SSD devices, or PCIe devices in the system, to improve system performance.

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 TieBreaker (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