# ScaleIO v1.32

## Lab Guide

August 2015

EMC²

# Copyright

Copyright © 1996, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014,2015 EMC Corporation. All Rights Reserved. EMC 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." EMC CORPORATION 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 EMC software described in this publication requires an applicable software license.

EMC2, EMC, Data Domain, RSA, EMC Centera, EMC ControlCenter, EMC LifeLine, EMC OnCourse, EMC Proven, EMC Snap, EMC SourceOne, EMC Storage Administrator, Acartus, Access Logix, AdvantEdge, AlphaStor, ApplicationXtender, ArchiveXtender, Atmos, Authentica, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, Avamar, Captiva, Catalog Solution, C-Clip, Celerra, Celerra Replicator, Centera, CenterStage, CentraStar, ClaimPack, ClaimsEditor, CLARiiON, ClientPak, Codebook Correlation Technology, Common Information Model, Configuration Intelligence, Configuresoft, Connectrix, CopyCross, CopyPoint, Dantz, DatabaseXtender, Direct Matrix Architecture, DiskXtender, DiskXtender 2000, Document Sciences, Documentum, eInput, E-Lab, EmailXaminer, EmailXtender, Enginuity, eRoom, Event Explorer, FarPoint, FirstPass, FLARE, FormWare, Geosynchrony, Global File Virtualization, Graphic Visualization, Greenplum, HighRoad, HomeBase, InfoMover, Infoscape, Infra, InputAccel, InputAccel Express, Invista, Ionix, ISIS, Max Retriever, MediaStor, MirrorView, Navisphere, NetWorker, nLayers, OnAlert, OpenScale, PixTools, Powerlink, PowerPath, PowerSnap, QuickScan, Rainfinity, RepliCare, RepliStor, ResourcePak, Retrospect, RSA, the RSA logo, SafeLine, SAN Advisor, SAN Copy, SAN Manager, Smarts, SnapImage, SnapSure, SnapView, SRDF, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix, Symmetrix DMX, Symmetrix VMAX, TimeFinder, UltraFlex, UltraPoint, UltraScale, Unisphere, VMAX, Vblock, Viewlets, Virtual Matrix, Virtual Matrix Architecture, Virtual Provisioning, VisualSAN, VisualSRM, Voyence, VPLEX, VSAM-Assist, WebXtender, xPression, xPresso, YottaYotta, the EMC logo, and where information lives, are registered trademarks or trademarks of EMC Corporation in the United States and other countries.

All other trademarks used herein are the property of their respective owners.

© Copyright 2014 EMC Corporation. All rights reserved. Published in the USA.

# Table of Contents


**LAB EXERCISE 1: CONNECTING TO THE LAB ENVIRONMENT ..... 3**

LAB 1: TASK 1 – CONNECT TO LAB ENVIRONMENT ..... 4

**LAB EXERCISE 2-A: SCALEIO LINUX/WINDOWS INSTALLATION ..... 8**

LAB 2-A: TASK 1 - INSTALL SCALEIO INSTALLATION MANAGER ..... 9
LAB 2-A: TASK 2 - CONNECT TO THE SCALEIO INSTALLATION MANAGER ..... 10
LAB 2-A: TASK 3 - UPLOAD THE INSTALLATION PACKAGES ..... 11
LAB 2-A: TASK 4 - UPLOAD THE CSV FILE ..... 14
LAB 2-A: TASK 6 - ADD ADDITIONAL SDS USING INSTALLATION MANAGER ..... 18

**LAB EXERCISE 2-B: SCALEIO VMWARE INSTALLATION ..... 22**

LAB 2-B: TASK 1 - INSTALL SCALEIO PLUG-IN FOR VSPHERE ..... 22
LAB 2-B: TASK 2 - INSTALL SVM TEMPLATE ..... 26
LAB 2-B: TASK 3 - INSTALL SDC INTO THE KERNEL OF THE SELECTED ESXI SERVERS ..... 28
LAB 2-B: TASK 4 – ADJUSTMENTS FOR LAB ENVIRONMENT ..... 31
LAB 2-B: TASK 5 – DEPLOY SCALEIO ONTO VMWARE ..... 35

**LAB EXERCISE 3-A: SCALEIO CONFIGURATION AND MANAGEMENT ..... 51**

LAB 3-A: TASK 1 – GUI INTRODUCTION ..... 51
LAB 3-A: TASK 2 – SCLI ..... 55
LAB 3-A: TASK 3 – CREATE VOLUMES ..... 59
LAB 3-A: TASK 4 – MAP VOLUMES TO SDC ..... 62
LAB 3-A: TASK 4 – MAP VOLUMES TO WINDOWS SDC ..... 67
LAB 3-A: TASK 5 – SNAPSHOT VOLUMES ..... 70
LAB 3-A: TASK 6 – CONSISTENCY GROUPS ..... 72
LAB 3-A: TASK 7 – INACTIVATION OF A PROTECTION DOMAIN ..... 77
LAB 3-A: TASK 8 – RENAME THE PROTECTION DOMAIN ..... 79
LAB 3-A: TASK 9 - RENAME STORAGE POOL ..... 80
LAB 3-A: TASK 9 - RENAME STORAGE POOL ..... 82

**LAB EXERCISE 3-B: SCALEIO CONFIGURATION AND MANAGEMENT - ESXI CLUSTER ..... 84**

LAB 3-B: TASK 1 - EXPLORE VSPHERE PLUGIN ..... 85
LAB 3-B: TASK 2 - CREATE AND MAP SCALEIO VOLUMES ..... 86
LAB 3-B: TASK 3 - CREATE DATASTORES ..... 90
LAB 3-B: TASK 4 - MIGRATE LIVE VM TO SCALEIO DATASTORE ..... 93
LAB 3-B: TASK 5 - MIGRATE LIVE VM TO OTHER HOST ..... 96

**LAB EXERCISE 4: MONITORING AND PERFORMANCE TUNING ..... 100**

LAB 4: TASK 1 - MONITORING, REBALANCE AND REBUILD WHEN ADDING/REMOVING AN SDS ..... 101
LAB 4: TASK 2 - ADJUSTING REBUILD AND REBALANCE POLICY ..... 109
LAB 4: TASK 3 - CONFIGURING QOS FOR A VOLUME ..... 112
LAB 4: TASK 4 - RAM CACHE DETERMINATION AND IMPLEMENTATION ..... 115
LAB 4: TASK 5 - MONITORING SCALEIO-REPORTED EVENTS ..... 119
LAB 4: TASK 6 - SET UP CENTRALIZED RSYSLOG SERVER FOR SCALEIO EVENT LOGGING ..... 122

**LAB EXERCISE 5-A: EXERCISING THE SCALEIO REST API ..... 128**

LAB 5-A: TASK 1 - VERIFY REST GATEWAY CONNECTION TO THE SCALEIO INSTANCE ..... 129
LAB 5-A: TASK 2 - PREPARE THE FIREFOX REST CLIENT..... 130
LAB 5-A: TASK 3 - ISSUE "GET" AND "POST" REQUESTS TO SCALEIO ..... 132
LAB X: END OF LAB..... 138

## Lab Exercise 1: Connecting to the lab environment

| | |
|---|---|
| **Purpose:** | Logon using your instructor-provided credentials to access the lab environment |

| | |
|---|---|
| **Tasks:** | ✓ Log in to the vdc<br>✓ Login to the vCloud Director<br>✓ Copy vApp from catalog to workspace<br>✓ Start vApp<br>✓ Work within VM desktop |

| | |
|---|---|
| **References:** | ScaleIO Lab Diagram (graphically showing relevant details for all VMs within a student pod – refer to Page 2) |

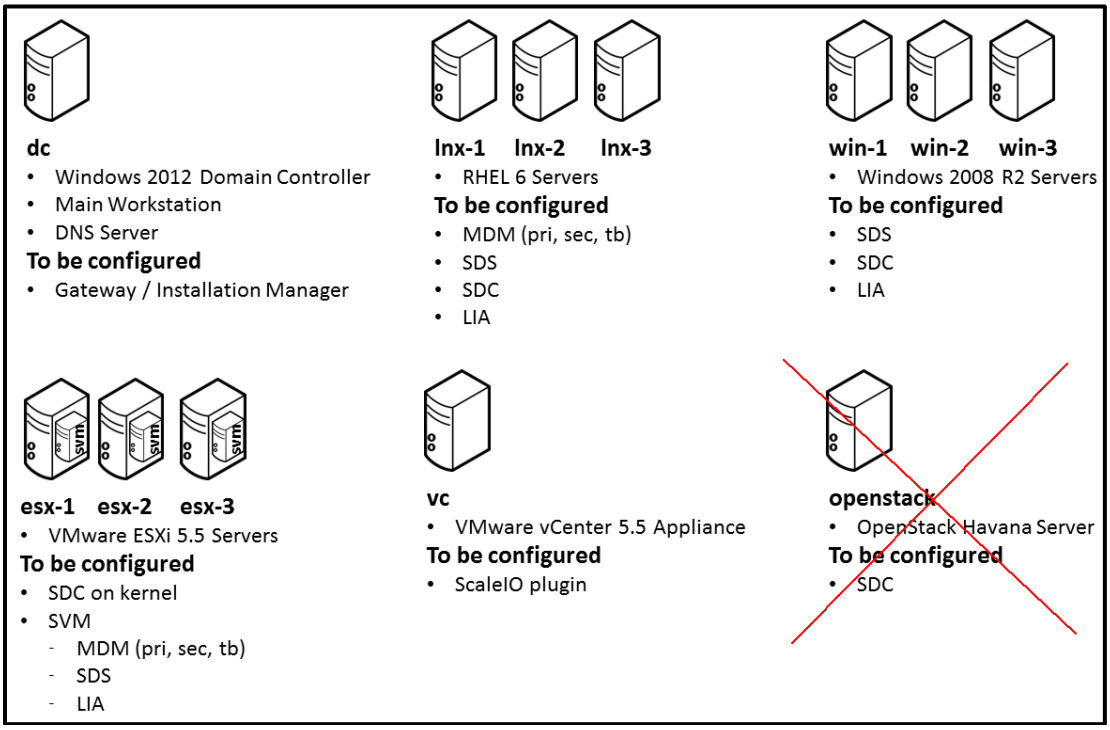

## Lab 1: Task 1 – Connect to lab environment

| Step | Action |
|---|---|
| | |
| 1. | Open your browser to https://vdc.emc.com<br> |
| 2. | Logon with the instructor provided credentials<br> 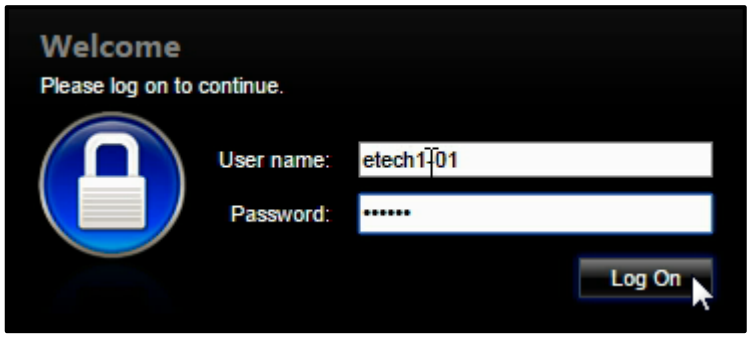 |
| 3. | Click on the link to the vCloud "Emerging Technology" datacenter (as shown)<br> |
| 4. | An Internet Explorer window will appear (this is being presented by a remote Citrix server, not the local client). If you receive a warning about a plugin, Allow it<br> |

| Step | Action |
|---|---|
| 5. | Log on with the same credentials as previously.<br> <br> |
| 6. | Click Add vApp from catalog<br> 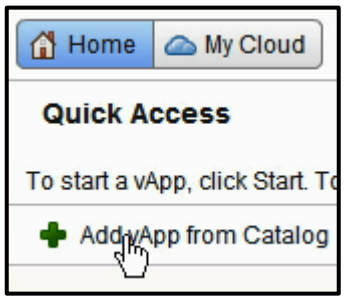 <br> |
| 7. | Select the ScaleIO I&M vApp then click Next<br> 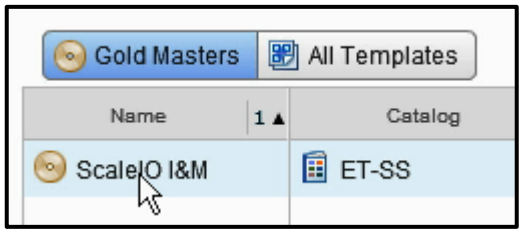 <br> |
| 8. | Let it select the default name and click Finish<br> <br> |

| Step | Action |
|---|---|
| 9. | Wait for the vApp to go from Creating to Stopped, then click the Start button.<br> 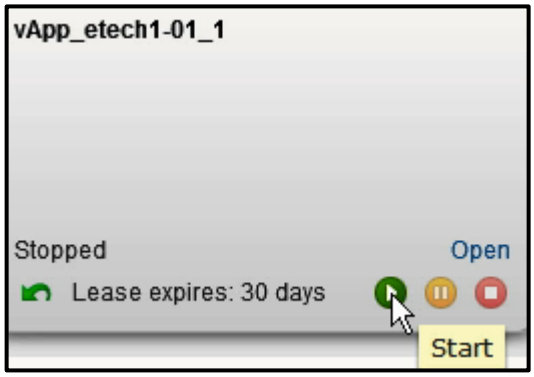 |
| 10. | Wait for it to go from Starting to Running, then click Open<br> |
| 11. | The “dc” is the only system you should directly interact with from here. Click inside the desktop to access the console<br> 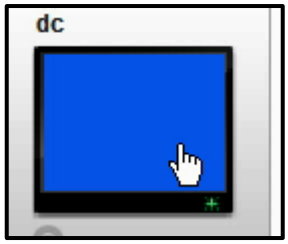 |

| Step | Action |
|---|---|
| 12. | Click inside the console to interact. Shown below are the functions of the icons.<br> 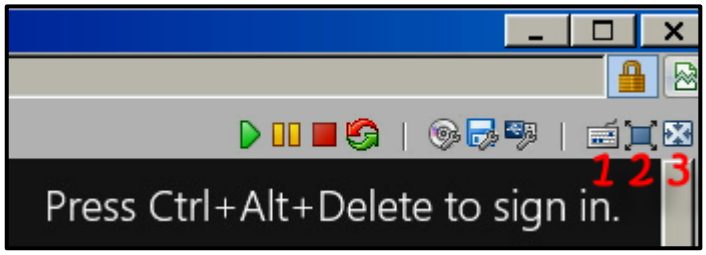 <br><br>1. Send ctrl-alt-del to console<br>2. Expand console to fullscreen view<br>3. Expand window to match console desktop resolution |

## Lab Exercise 2-A: ScaleIO Linux/Windows Installation

| | |
|---|---|
| **Purpose:** | Use the ScaleIO V Installation Manager (IM) graphical utility to install and expand a ScaleIO cluster on physical platforms running a mix of Linux and Windows. |

| | |
|---|---|
| **Tasks:** | ✓ Connect to the dc desktop<br>✓ Install and configure the ScaleIO V Installation Manager<br>✓ Install a three-node ScaleIO cluster of Windows/Linux hosts, using the Installation Manager GUI<br>✓ Expand the cluster by adding one more Windows/Linux host as an SDS, also using Installation Manager |

| | |
|---|---|
| **References:** | ScaleIO Lab Diagram (graphically showing relevant details for all VMs within a student pod – refer to Page 2) |

## Lab 2-A: Task 1 - Install ScaleIO Installation Manager

| Step | Action |
|---|---|
| | In this section, you will install and run the Installation Manager onto the DC system |
| 13. 1 | Log into the DC system. This is your access point for all other servers assigned to your pod. |
| 14. 2 | Open Windows Explorer to the following folder:<br>C:\share\sw\Gateway_for_Windows<br> 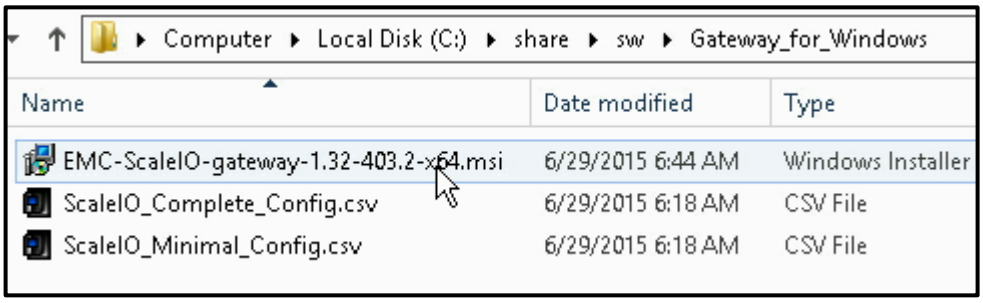 <br> |
| 15. 3 | Double-click (run) the **EMC-ScaleIO-Gateway….msi** file.<br><br>Proceed through the install, accepting the defaults and the license agreement.<br><br>Enter the new admin password that will be used to access IM: **P@ssw0rd**<br> 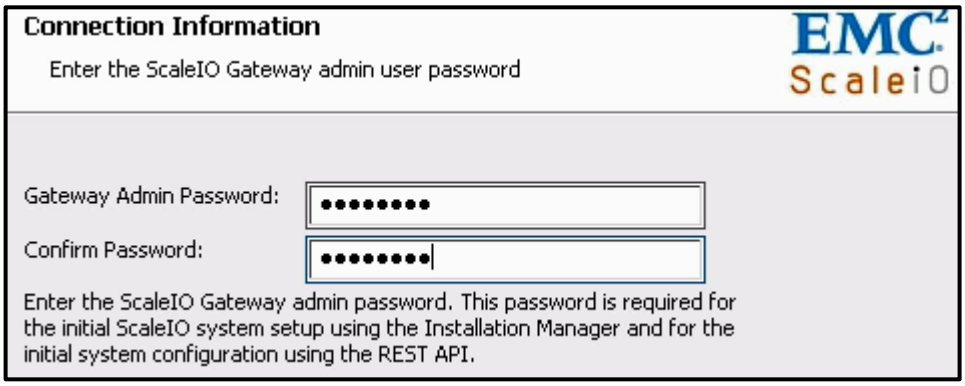 <br><br><br>Once finished, close the Explorer windows and continue. |

## Lab 2-A: Task 2 - Connect to the ScaleIO Installation Manager

| Step | Action |
|---|---|
| 1. | ***<u>From the dc desktop</u>***, use Firefox to access the URL for IM: https://dc.emc.edu 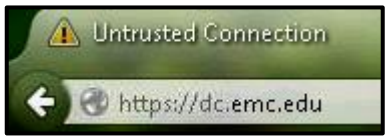  |
| 2. | Accept the risks. Add the exception. 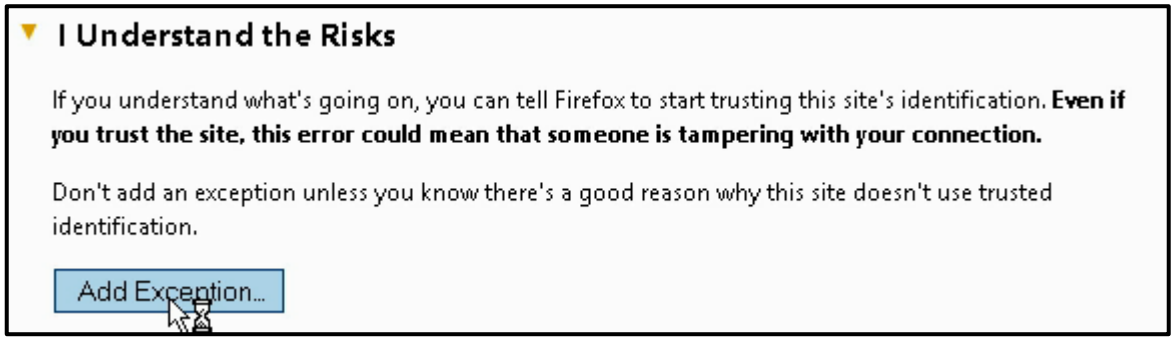  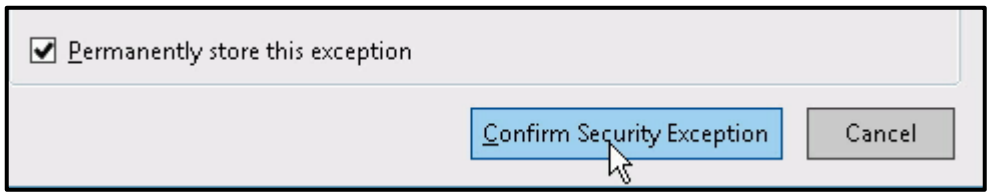  |

| Step | Action |
|---|---|
| 3. | Logon as User: admin \| Password: P@ssw0rd<br> 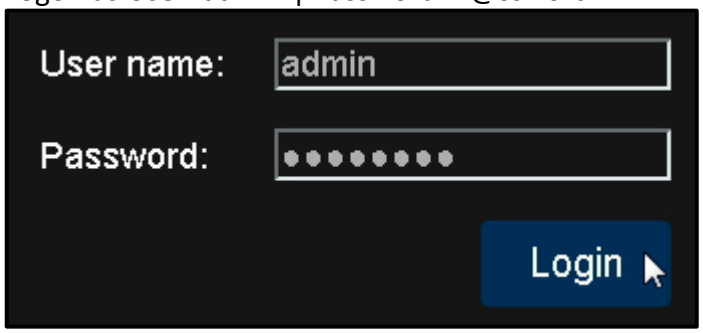 <br><br>It is recommended to allow Firefox to save these credentials for convenience<br> 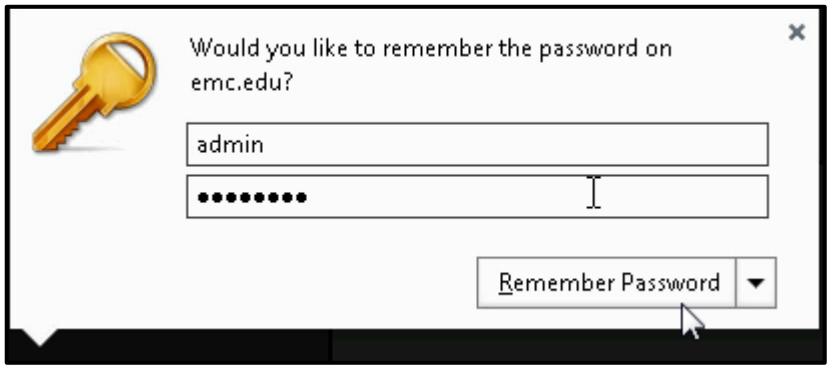 <br> |

## Lab 2-A: Task 3 - Upload the Installation Packages

| Step | Action |
|---|---|
| | Before any installation can take place, the ScaleIO packages need to be uploaded to the installation manager gateway system. |

| Step | Action |
|---|---|
| 1. | Along the top of the IM screen, click **Packages**.<br> <br><br>The **Manage Installation Packages** window appears. Click **Browse**<br> 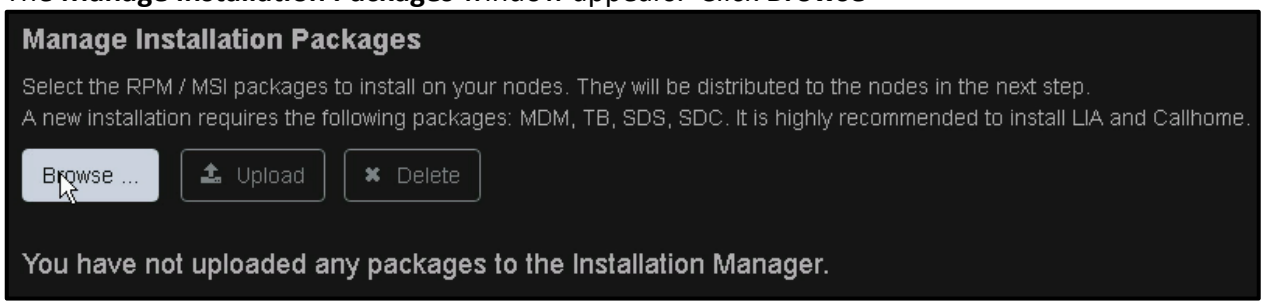 <br> |
| 2. | In the dialog box, open the **c:\share\sw\Windows** folder.<br>Select all six MSI files found in that folder and click **Open**.<br> 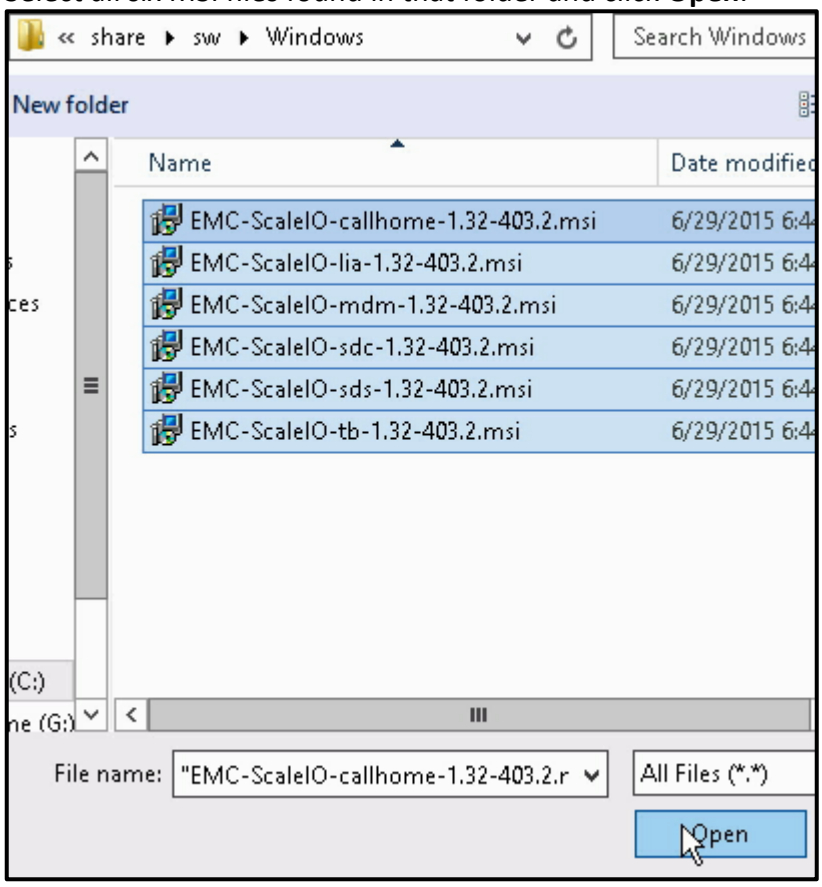 <br><br>**Hint:** Click the first file, then hold **Shift** and click the last file to select them all at once. |

| Step | Action |
|---|---|
| 3. | Click **Upload**.<br> 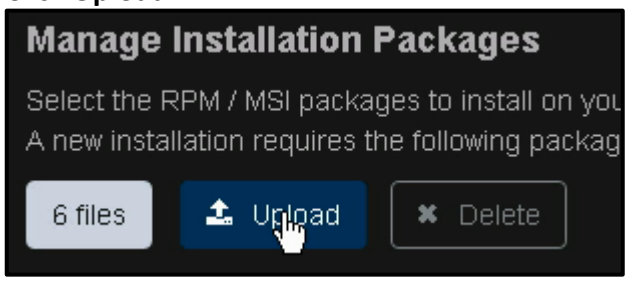 <br><br>The uploaded installation packages appear in the file table:<br> 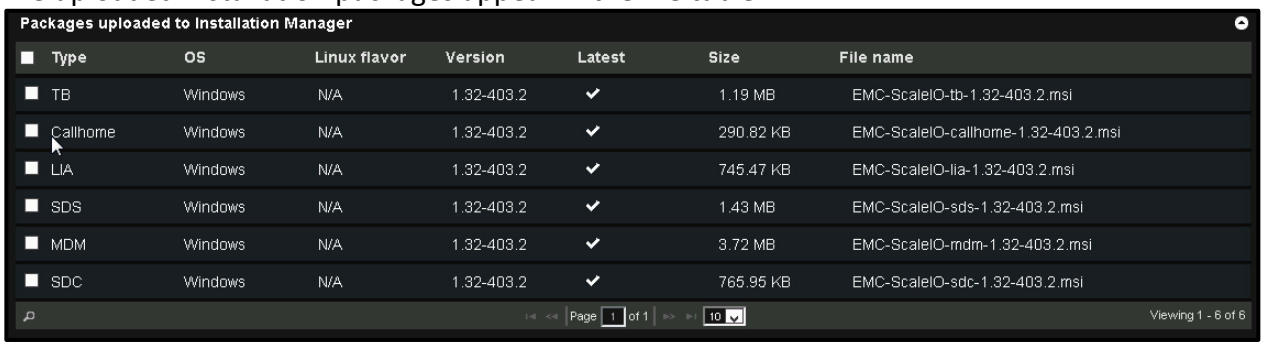 <br> |
| 4. | Repeat Steps 2 and 3 using the **c:\share\sw\RHEL6** folder, and upload the six RPMs found in that folder.<br> 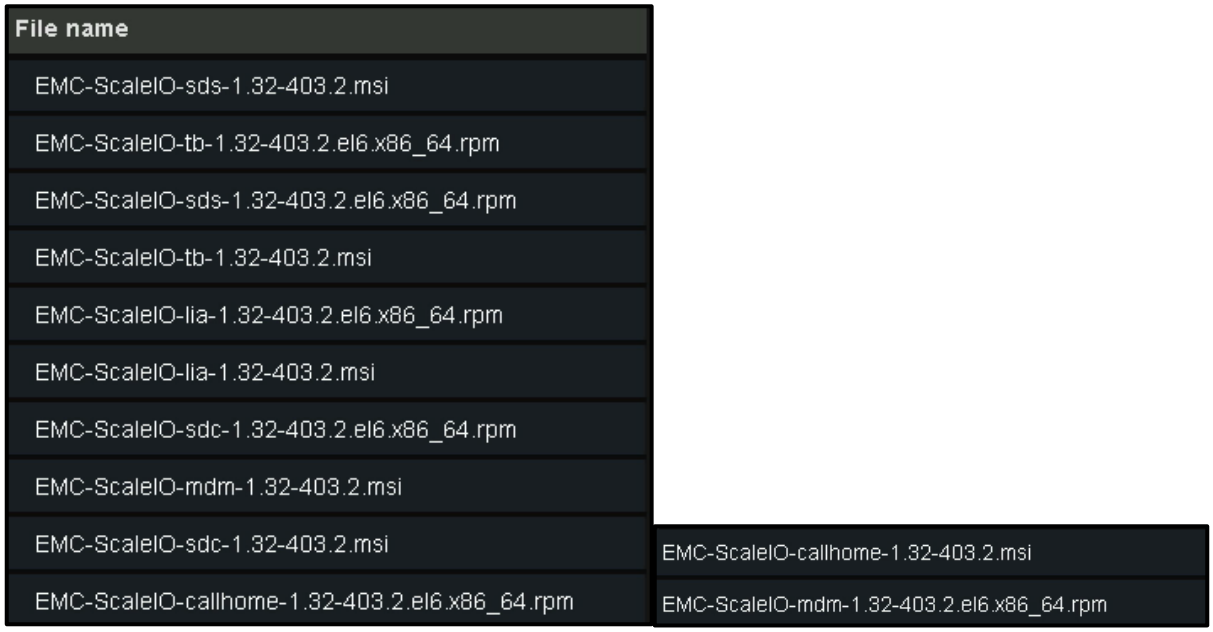 <br> |
| 5. | When you are finished uploading the files, click **Proceed to Install >** to continue.<br> <br> |

## Lab 2-A: Task 4 - Upload the CSV File

| Step | Action |
|---|---|
| | The CSV file provides the installation manager system with the data required to complete the installation. It includes the host information, along with roles, networking, devices etc. |
| 1. | The browser goes to the **Installation Topology** screen. Click the **Browse** button.: 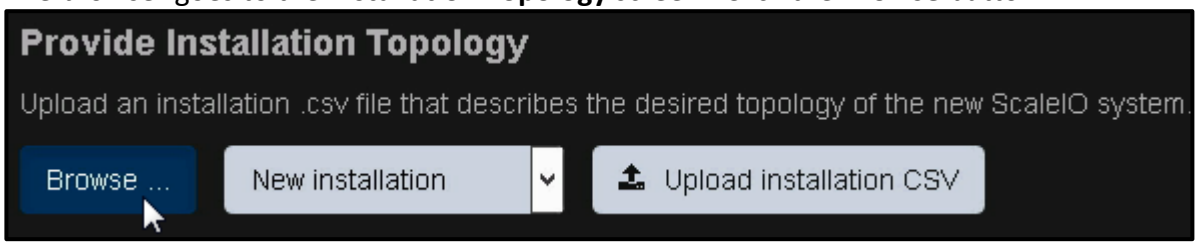  |
| 2. | Select the following file C:\share\misc\ **ScaleIO_Initial_Config.csv**.<br>Click **Open**. The file name appears in place of the Browse button. 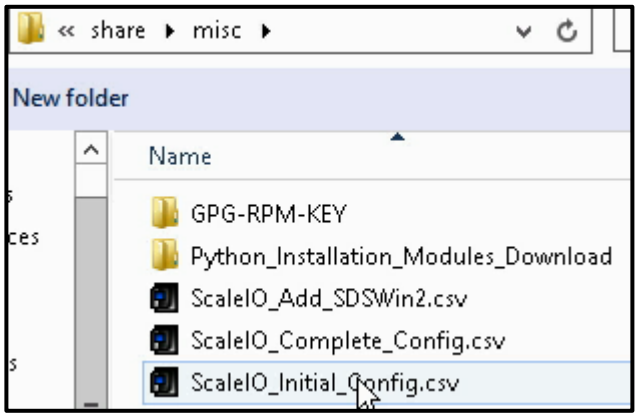  |
| 3. | For your reference, here is the file in Microsoft Excel: 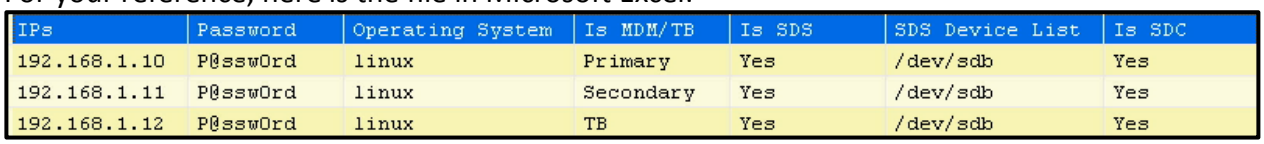 |

| IPs | Password | Operating System | Is MDM/TB | Is SDS | SDS Device List | Is SDC |
|---|---|---|---|---|---|---|
| 192.168.1.10 | P@sswOrd | linux | Primary | Yes | /dev/sdb | Yes |
| 192.168.1.11 | P@sswOrd | linux | Secondary | Yes | /dev/sdb | Yes |
| 192.168.1.12 | P@sswOrd | linux | TB | Yes | /dev/sdb | Yes |

**Note:** For your convenience, this CSV file has been pre-populated with all the relevant data for your specific ScaleIO hosts. In a real deployment, you will need to manually edit the configuration file suitably, before uploading it to the Installation Manager.

| | |
|---|---|
| 4. | Click the **Upload installation CSV** button. |

| Step | Action |
|---|---|
| 5. | At the top of the installation configuration screen, set the MDM and LIA passwords to P@ssw0rd. Accept the license Agreement, and remove the check for Call Home (which we will not be configuring during this lab exercise.<br> 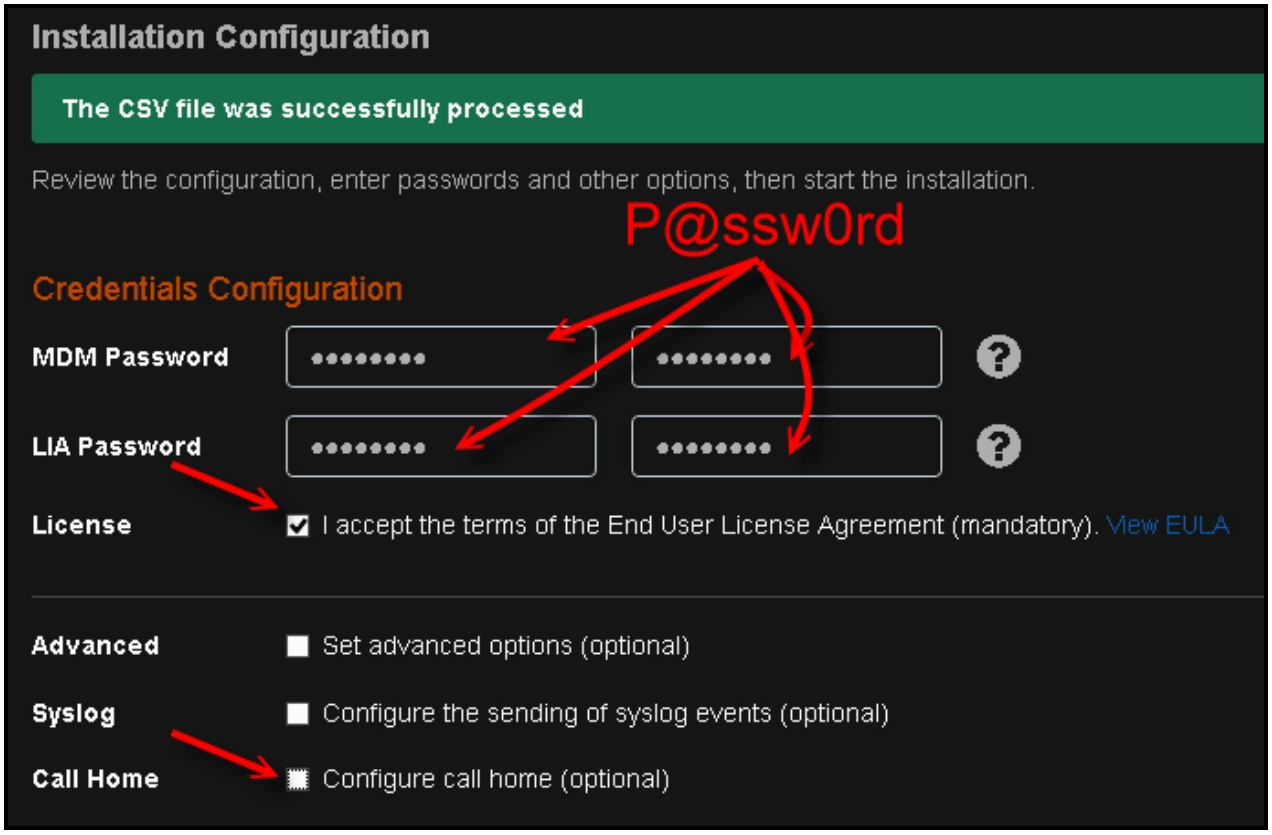 <br> |

| Step | Action |
|---|---|
| 6. | Confirm the topology (but don't change anything).<br>Topology<br>MDMs and Tie-Breaker<br>Role / MDM IPS(s) / OS / User name / Password<br>Primary / 192.168.1.10 / Linux / root / P@ssw0rd<br>Secondary / 192.168.1.11 / Linux / root / P@ssw0rd<br>Tie-Breaker / 192.168.1.12 / Linux / root / P@ssw0rd<br>SDS List<br>IP / All IPs / OS / PD / SP / Devices / User name / Password<br>192.168.1.12 / 192.168.1.12 / Linux / default / default / /dev/sdb / root / show<br>192.168.1.11 / 192.168.1.11 / Linux / default / default / /dev/sdb / root / show<br>192.168.1.10 / 192.168.1.10 / Linux / default / default / /dev/sdb / root / show<br>Page 1 of 1 — Viewing 1 - 3 of 3<br>SDC List<br>IP / OS / User name / Password<br>192.168.1.12 / Linux / root / show<br>192.168.1.11 / Linux / root / show<br>192.168.1.10 / Linux / root / show<br>Page 1 of 1 — Viewing 1 - 3 of 3<br>Click<br>Start Installation > |
| 7. | Click on either monitor screen to bring up the installation monitor<br>ScaleIO Installer<br>Home / Packages / Install / Monitor<br>Installation started successfully.<br>The next step is to navigate to the monitor tab where you will monitor and approve the installation.<br>Starting the query phase of the installation ... |

| Step | Action |
|---|---|
| 8. | The installation performs the following phases: upload, install, and configure.<br><br>**Note:** If you get an error message during the process, you can retry the failed operations.<br> 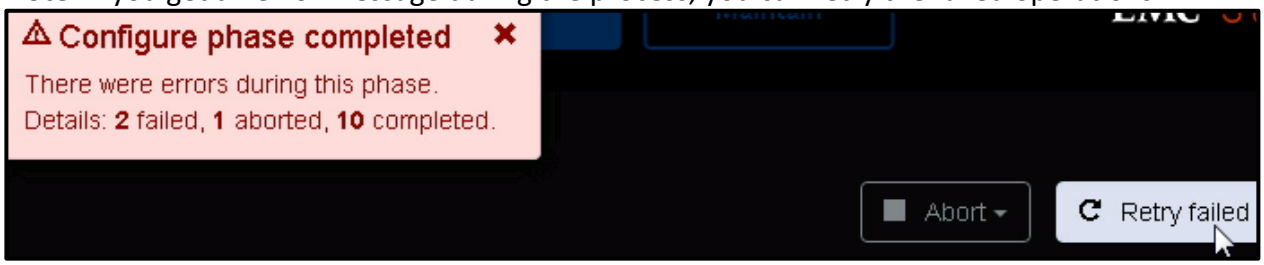 |
| 9. | The Query phase starts automatically. When it finishes without error, close the "completed" box, and click Start upload phase<br> 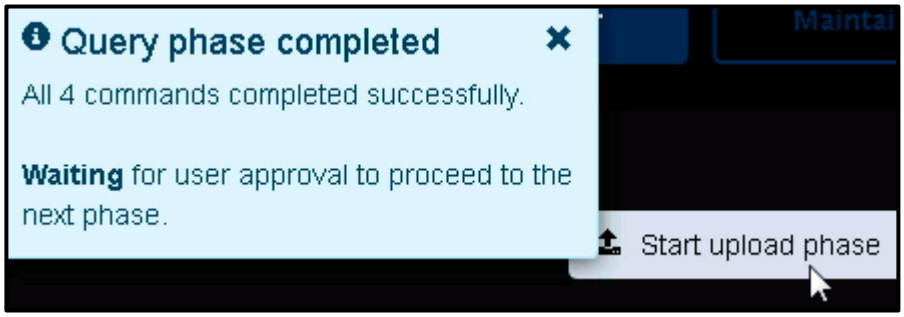 |
| 10. | When the Upload phase finishes without error, close the "completed" box, and click Start install phase<br> 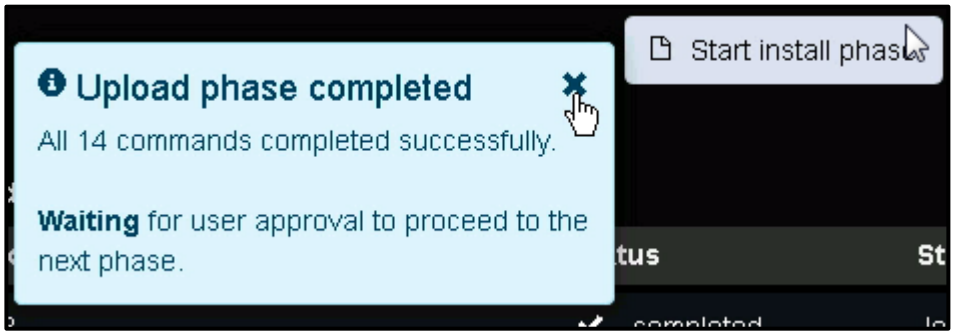 |

| Step | Action |
|---|---|
| 11. | When the Install phase finishes without error, close the “completed” box, and click Start Configure phase<br> 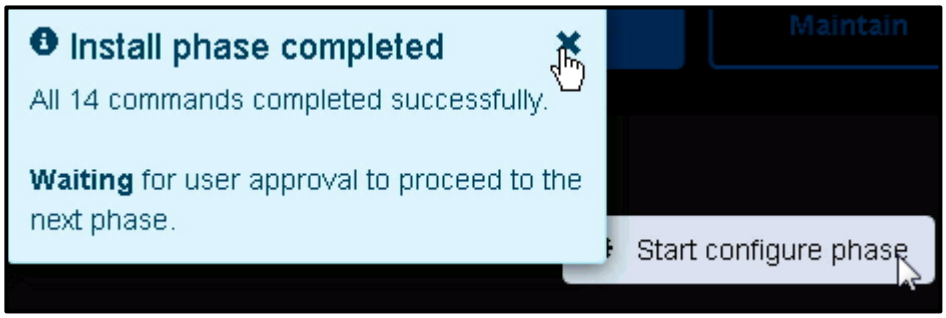 <br> |
| 12. | Configure is the final phase. When it finished, the installation and configuration is complete on all nodes. Click Mark Operation Completed.<br> 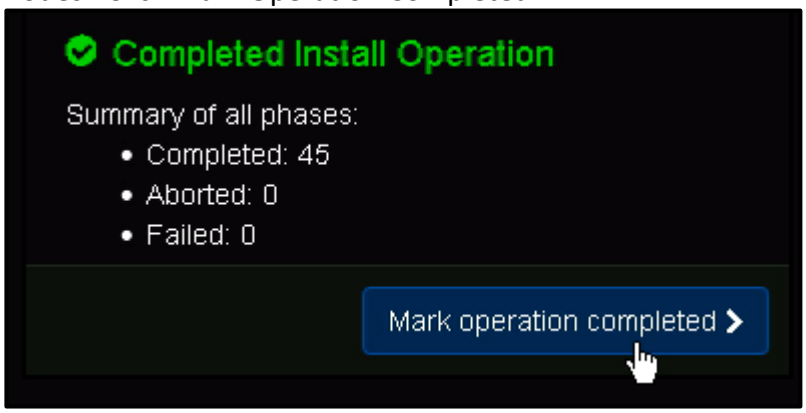 <br> |

## Lab 2-A: Task 6 - Add additional SDS using Installation Manager

| Step | Action |
|---|---|
| | In this section you will use a csv file to add a single SDS system to the cluster you’ve just constructed |
| 1. | Within the installation GUI, Click on the Install box<br> 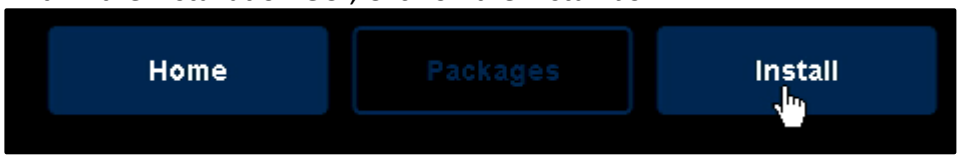 <br> |

| Step | Action |
|---|---|
| 2. | In the provide installation topology screen, click Add to existing system in the middle, then click Browse<br> <br> |
| 3. | Browse to c:\share\misc, and select **ScaleIO_Add_SDSWin2.csv** file<br> 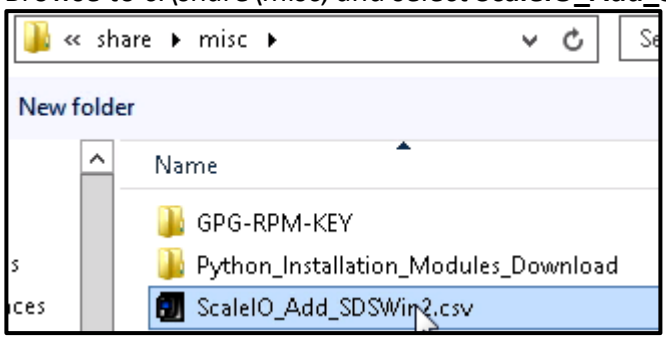 <br><br>Here are the contents of the file<br>Domain \| Username \| Password \| Operating System \| Is MDM/TB \| MDM Mgmt IP<br>localhost \| administrator \| P@sswOrd \| windows \| No<br>root \| P@sswOrd \| linux \| Primary<br>root \| P@sswOrd \| linux \| Secondary<br>root \| P@sswOrd \| linux \| TB<br>MDM IPs \| Is SDS \| SDS Name \| SDS All IPs \| SDS-SDS Only IPs \| SDS-SDC Only IPs<br>Yes \| SDSWin2 \| 192.168.1.21<br>192.168.1.10<br>192.168.1.11<br>192.168.1.12<br>Protection Domain \| Fault Set \| SDS Device List \| SDS Pool List \| Optimize IOPS \| Is SDC<br>default \| e \| default \| No \| No<br><br>Click Upload |

| Step | Action |
|---|---|
| 4. | On the **Review Installation Configuration** screen:<br>Enter admin \| P@ssw0rd for MDM credentials, and P@ssw0rd for the LIA password.<br>Accept the license.<br>Do not configure call home.<br> 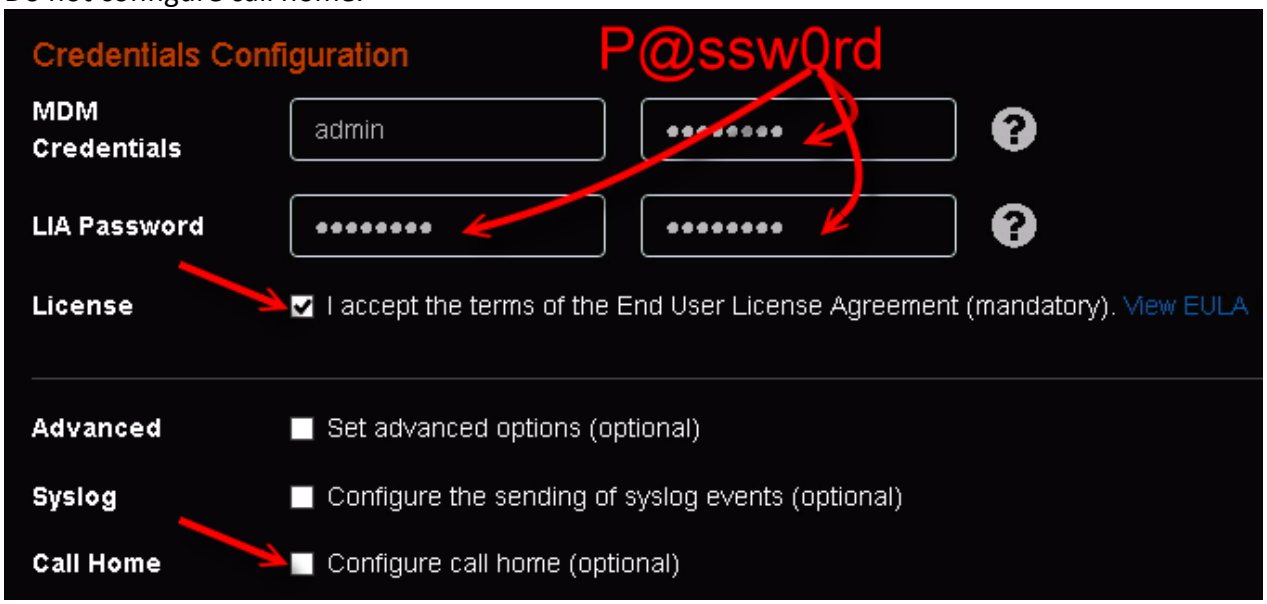 <br> |
| 5. | Take a moment to review the displayed information regarding the new SDS.<br> 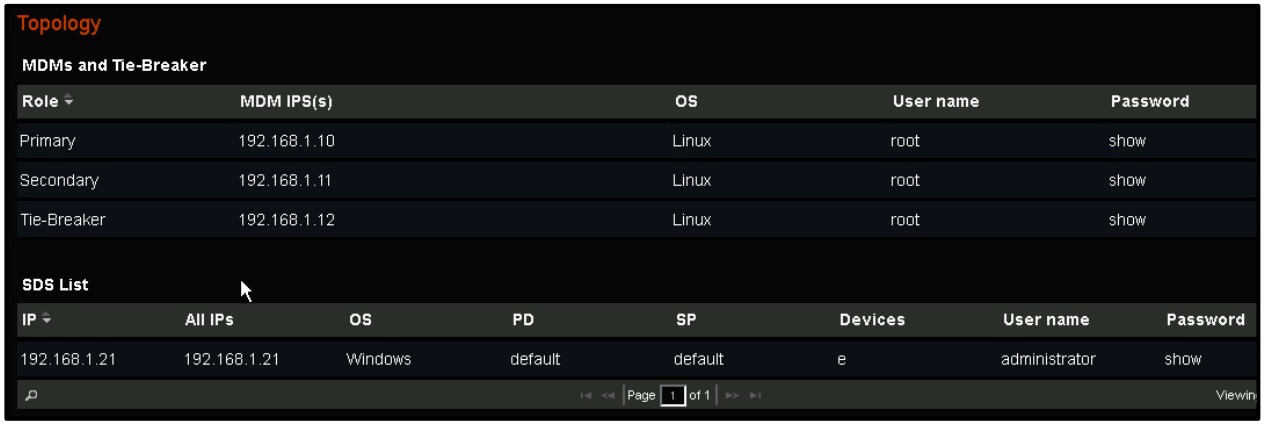 <br> |
| 6. | Click Start Installation<br> <br> |
| 7. | Click Monitor<br> <br> |

| Step | Action |
|---|---|
| 8. | The phases will be identical to the previous phases. If you see an error or aborted operation, click Retry Failed. Otherwise, follow the instructions similar to the previous section, |
| 9. | When the configure phase completes, click mark operation completed<br> 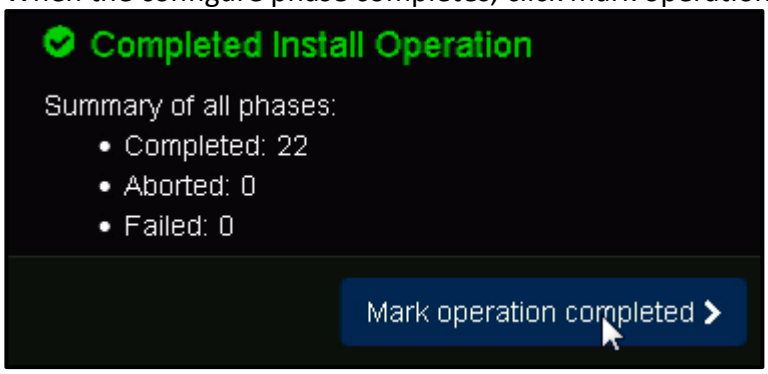 <br> |
| 10. | Minimize the browser and Click on the "lnx-1" icon<br>*(it will automatically connect to lnx-1 via putty and logon as root / P@ssw0rd)*<br> <br> |
| 11. | Run the following commands:<br>**`scli --login --username admin --password P@ssw0rd`**<br>**`scli --query_cluster`**<br>`Using username "root".`<br>`Last login: Mon Jul 27 11:50:54 2015 from 192.168.1.50`<br>`[root@Lnx-1 ~]# scli --login --username admin --password P@ssw0rd`<br>`Logged in. User role is Administrator. System ID is 3ce598964670c629`<br>`[root@Lnx-1 ~]# scli --query_cluster`<br>`Mode: Cluster, Cluster State: Normal, Tie-Breaker State: Connected`<br>`Primary MDM IP: 192.168.1.10`<br>`Secondary MDM IP: 192.168.1.11`<br>`Tie-Breaker IP: 192.168.1.12`<br>`Management IP: 192.168.1.10, 192.168.1.11`<br>`Name: N/A`<br>`[root@Lnx-1 ~]#` |

*End of Lab Exercise 2-A*

## Lab Exercise 2-B: ScaleIO VMware Installation

| | |
|---|---|
| **Purpose:** | Use the ScaleIO vSphere Web Client plug-in to install and expand a ScaleIO cluster on the VMware ESXi platform. |

| | |
|---|---|
| **Tasks:** | ✓ Install the ScaleIO plug-in<br>✓ Install a three-node ScaleIO cluster of VMware ESX hosts using the ScaleIO Deployment Wizard<br>✓ Assess validity and completeness of the installed configuration, including: executing basic ScaleIO commands on one of the deployed SVM appliances |

## Lab 2-B: Task 1 - Install ScaleIO Plug-in for vSphere

| Step | Action |
|---|---|
| | This task has you installing the ScaleIO plugin into the vCenter server. It is handled via a powercli script. |
| 1. | Click on the PowerCLI icon on the desktop<br> |
| 2. | Open an explorer window to c:\share\sw\VMware\Plugin<br> 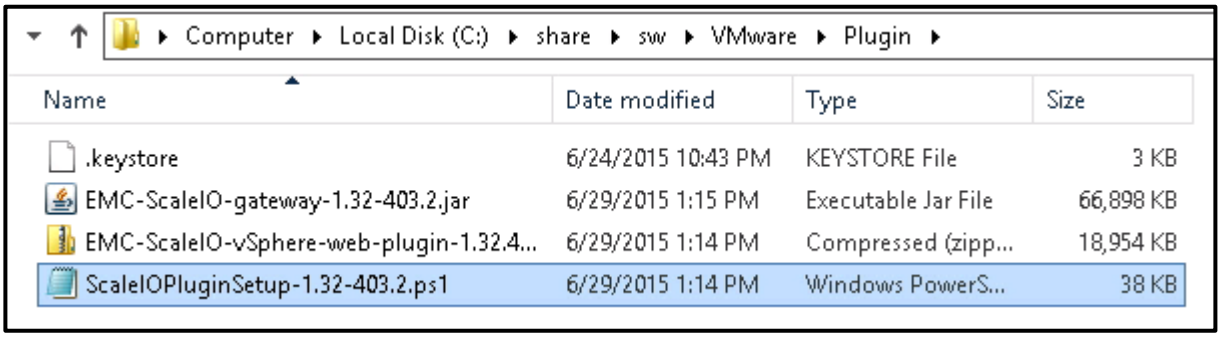 |

| Step | Action |
|---|---|
| 3. | Drag the ScaleIO…ps1 file into the PowerCLI window<br> 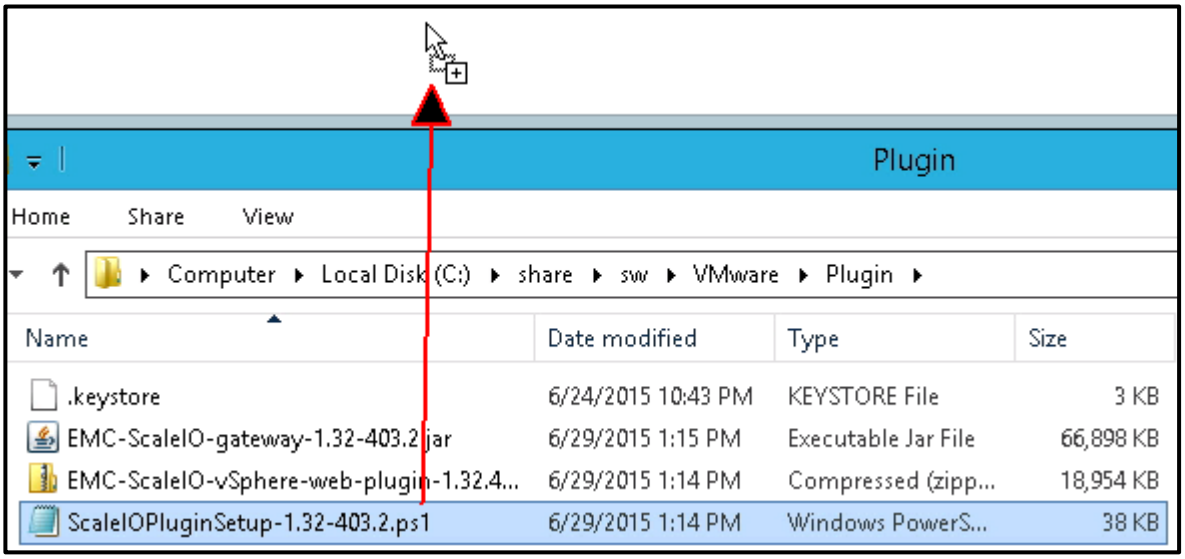 <br><br>`PowerCLI C:\Program Files (x86)\VMware\Infrastructure\vSphere PowerCLI> C:\share\sw\VMware\Plugin\ScaleIOPluginSetup-1.3`<br>`2-403.2.ps1` |
| 4. | Hit enter to launch the ps1 file. Enter the following information at the prompts:<br><br>- vCenter vc.emc.edu<br>- Username root<br>- Password P@ssw0rd<br><br>`PowerCLI C:\Program Files (x86)\`<br>`2-403.2.ps1`<br><br>`cmdlet ScaleIOPluginSetup-1.32-4`<br>`Supply values for the following`<br>`(Type !? for Help.)`<br>`vCenter: vc.emc.edu`<br>`Username: root`<br>`Password: ********` |
| 5. | Enter 1 to register the plugin, y to continue, and s for Standard registration<br><br>`Choose mode:`<br>`1 - Register ScaleIO plug-in`<br>`2 - Unregister ScaleIO plug-in`<br>`3 - Create SVM template`<br>`4 - Exit the script`<br>`: 1`<br><br>`Welcome to ScaleIO Plugin Register Utility`<br><br>`Important notice to users upgrading to ScaleIO release 1.32:`<br>`1. ESXi 5.1 is no longer supported. ESXi 5.5 and above are supported.`<br>`2. Running SDC within an SVM is no longer supported. SDC must be installed in the ESXi host.`<br>`If your configuration isn't supported, DO NOT UPGRADE.`<br>`Contact Support for additional info and further assistance.`<br>`Do you wish to continue? Y / N: y`<br>`Do you want to run the Standard or Advanced registration? [s\|a]: s`<br>`Waiting for internal server to load (attempt 1)...` |

| Step | Action |
|---|---|
| 6. | After several seconds you will prompted to log in to the vSphere client (do not hit enter yet)<br> 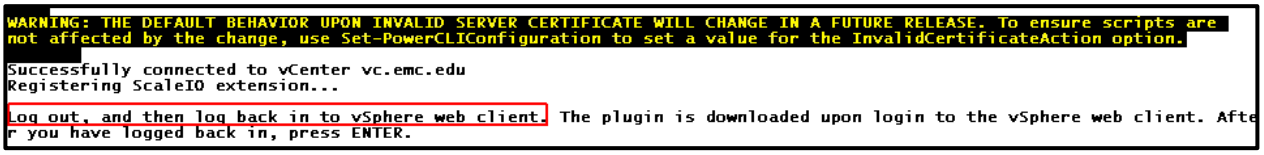 <br> |
| 7. | Connect to the vSphere web client by using the desktop icon (or by pointing firefox to https://vc.emc.edu:9443)<br> <br> |
| 8. | Logon as root \| P@ssw0rd<br> 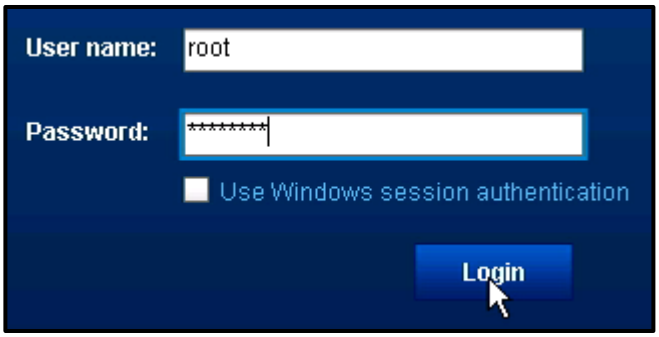 <br> |

| Step | Action |
|---|---|
| 9. | From the vSphere home screen, click on the Home tab – notice the EMC ScaleIO plugin (no need to click on it).<br> 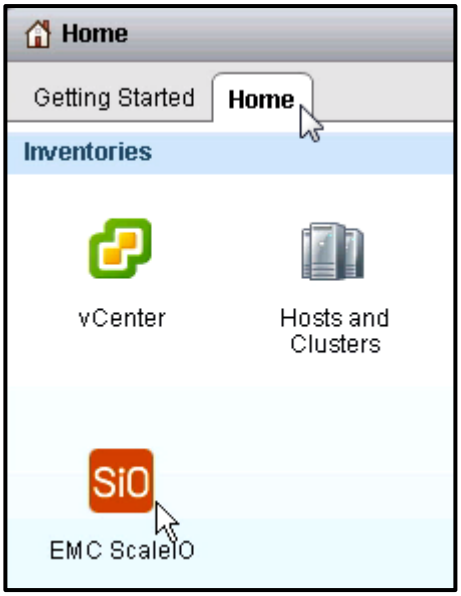 <br> |
| 10. | Return to PowerCLI and hit Enter to continue<br> 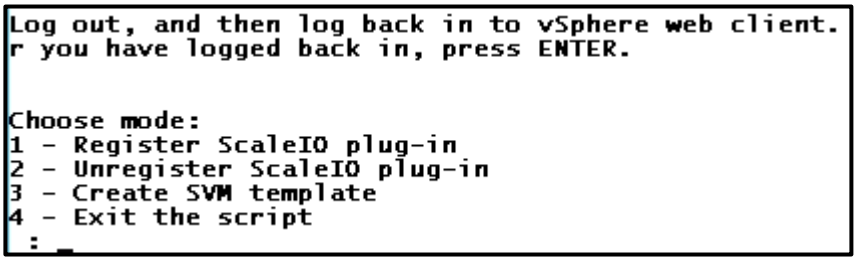 <br><br>Do not exit PowerCLI. |

## Lab 2-B: Task 2 - Install SVM Template

| Step | Action |
|---|---|
| | This task has you continuing to work within the ScaleIO Registration script, to create a ScaleIO Virtual Machine (SVM) template. This template will be later used by the ScaleIO deployment wizard to create virtual machines enabling MDM, SDS, LIA, and Gateway functionality.<br><br>You will need to assign it to a specified datacenter, specified datastore(s), and select the "ova" file to be used. |
| 1. | Enter three from the PowerCLI script menu to start the template registration<br><br>`Choose mode:`<br>`1 - Register ScaleIO plug-in`<br>`2 - Unregister ScaleIO plug-in`<br>`3 - Create SVM template`<br>`4 - Exit the script`<br>`: 3`<br>`Disconnecting from the vCenter vc.emc.edu...`<br><br>`Welcome to ScaleIO Plugin Create SVM Template Utility` |
| 2. | Enter Datacenter as the datacenter<br><br>`Welcome to ScaleIO Plugin Create SVM Template Utility`<br><br>`Enter a datacenter name: Datacenter`<br>`Datacenter is: Datacenter` |

| Step | Action |
|---|---|
| 3. | When prompted for the location of the ova file, open an explorer window and drag the following file into the powercli window, and hit enter:<br>C:\share\sw\VMware\ScaleIOVM…ova<br><br>*(or you can always just type it in as shown)*<br><br> 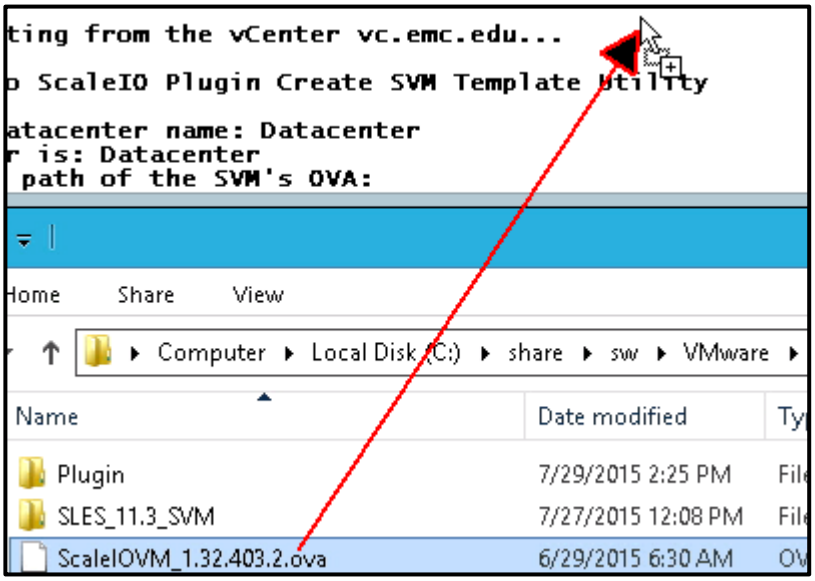 <br><br> <br> |
| 4. | When prompted for a datastore, enter DS30. When prompted for a second datastore, just hit enter (to skip).<br> 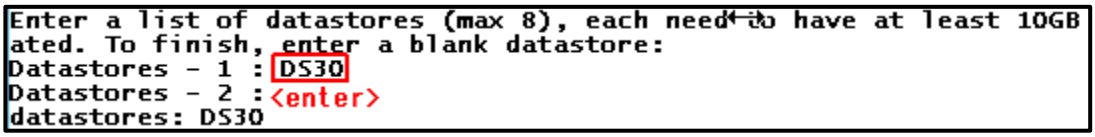 <br> |
| 5. | The process will begin immediately, and will take several minutes. When you are returned to the script menu, enter 4 to exit the script. Then type exit to exit PowerCLI<br><br> 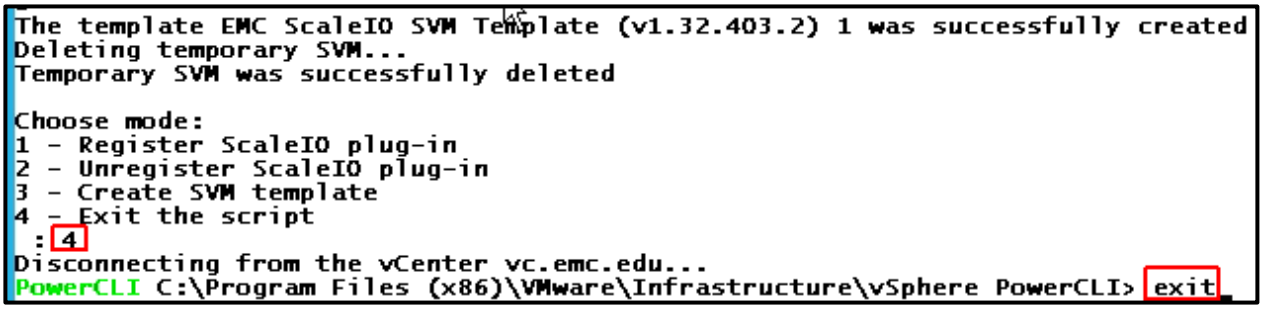 <br> |

## Lab 2-B: Task 3 - Install SDC into the kernel of the selected ESXi servers

| Step | Action |
|---|---|
| | In this task, you will install the SDC onto the kernels of the 3 ESXi hosts. It is done through an installer included with the ScaleIO plugin. It requires a reboot of the ESXi server. |
| 1. | In the vSphere Web Client (root \| P@ssw0rd), Home screen, click on the ScaleIO plugin 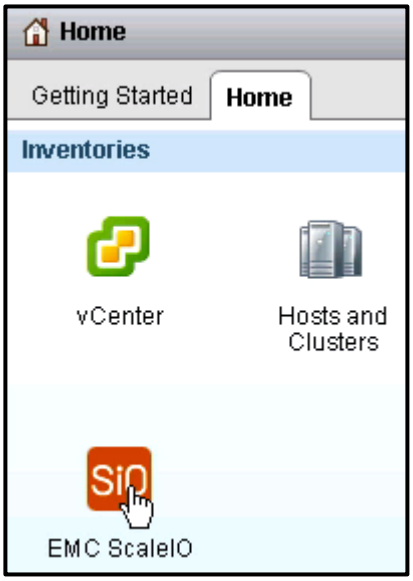  |
| 2. | Under Basic tasks, click on Install SDC on ESX |

| Step | Action |
|---|---|
| 3. | Select all three ESXi servers, enter P@ssw0rd as the password (click show passwords to help ensure it is entered correctly), and click Install<br> 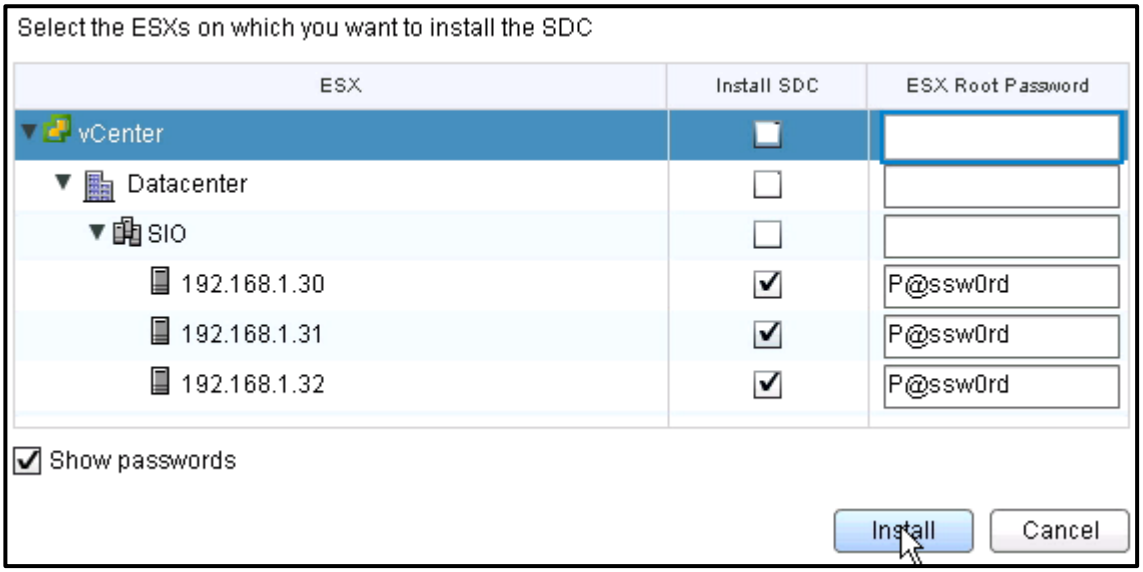 <br> |
| 4. | The installation will commence. When it is finished, click Finish<br> 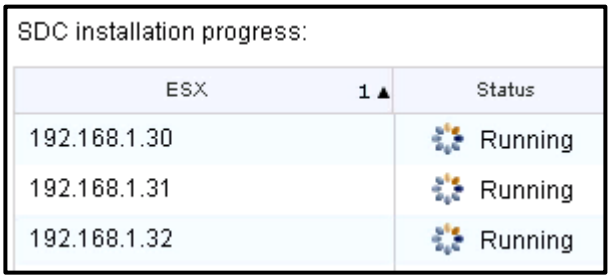 <br><br> 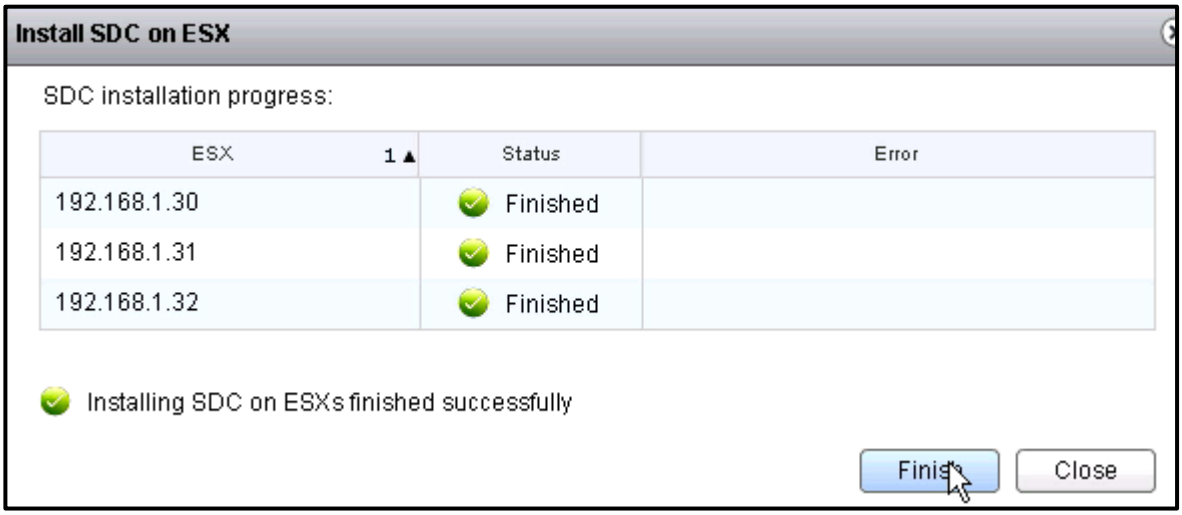 <br> |

| Step | Action |
|---|---|
| 5. | Click OK on the popup<br> <br> |
| 6. | On the left side of the vsphere home page, click on Hosts and Clusters<br> 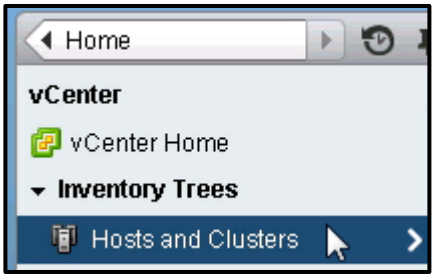 <br> |
| 7. | Right-click on the 192.168. ESX server, and click Reboot. Then click OK on the popup to reboot the host<br> 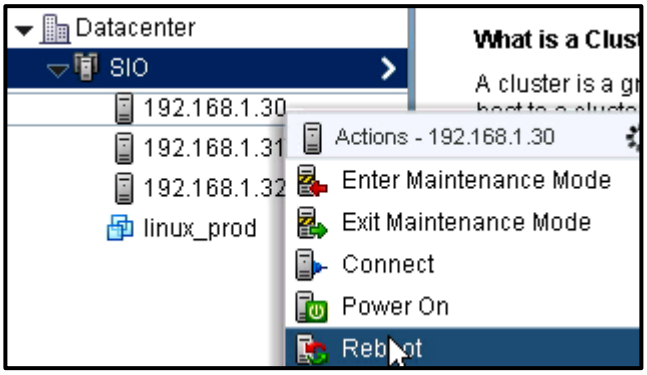 <br><br> <br> |

| Step | Action |
|---|---|
| 8. | Repeat previous step for the other two ESX servers. Wait until all three servers have rebooted before continuing. They will be unresponsive for a minute or two. Wait until they appear "normal" (bottom screen shot) before continuing. You can click the refresh button in the vSphere Web Client for a more accurate update<br> 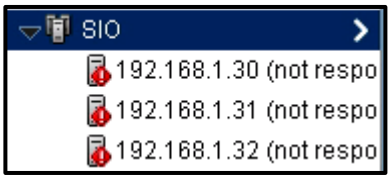 <br><br> 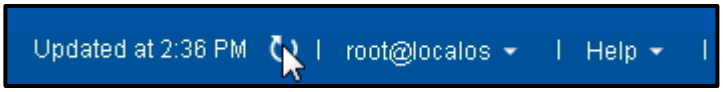 <br><br> 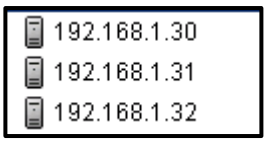 <br> |

## Lab 2-B: Task 4 – Adjustments for lab environment

| Step | Action |
|---|---|
| | This task has you making adjustments necessary to support the lab environment. These are not the recommended settings in a production environment.<br><br>In order to create a scalable lab allowing multiple concurrent connections, we are using Virtual ESXi servers. These allow for full ScaleIO functionality, but require minor settings changes.<br><br>One is to "allow vmdk creation". Normally, the selected SDS devices are passed to the SVM as raw devices (using VMware's RDM capability). This will not work here. This option formats the devices with VMFS, configured as a standard Datastore, and creates a single VMDK on each, which is then presented as a virtual disk to the SVM. Because the VMDK's use nearly all the available space on the VMFS, expect to see warnings in the vSphere client about datastore free space. You can ignore these warnings.<br><br>The other adjustment speeds up the deployment time. When VMDK's are created, the "eager zeroed" option is the default. This writes a "0" to all bits of the VMDK. While it creates a more efficient VMDK in production, the time to write this is excessive. Disabling this speeds up the process by 20 minutes or more. |

| Step | Action |
|---|---|
| 1. | Click on the home button and return to the ScaleIO plugin<br> 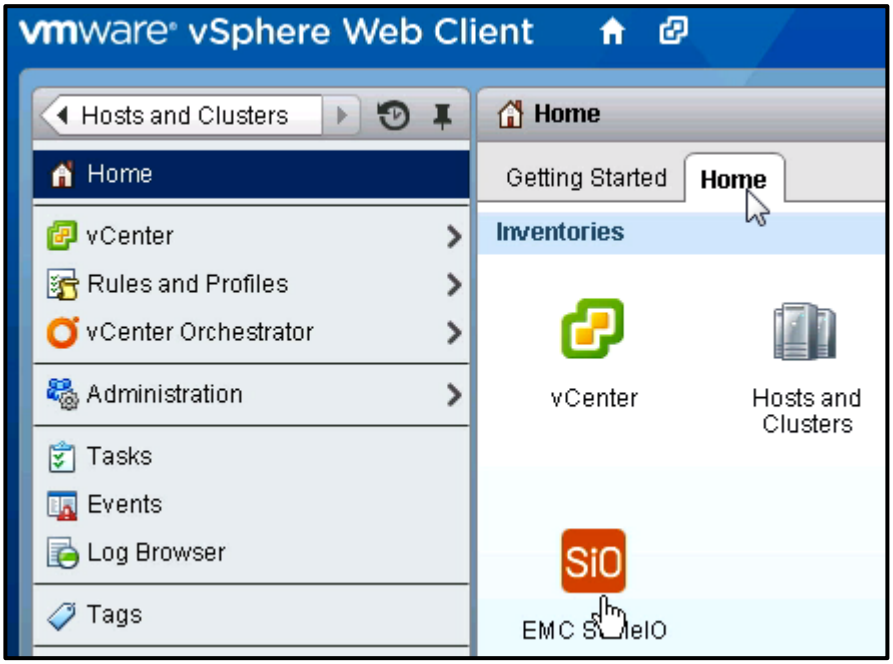 <br> |
| 2. | Click on Advanced Settings on the right<br> 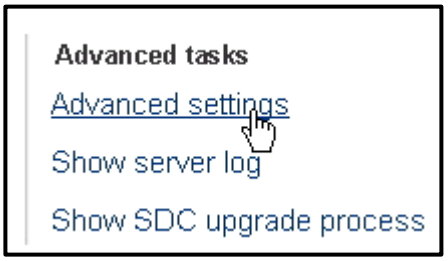 <br> |

| Step | Action |
|---|---|
| 3. | Check Enable VMDK creation and click OK<br> 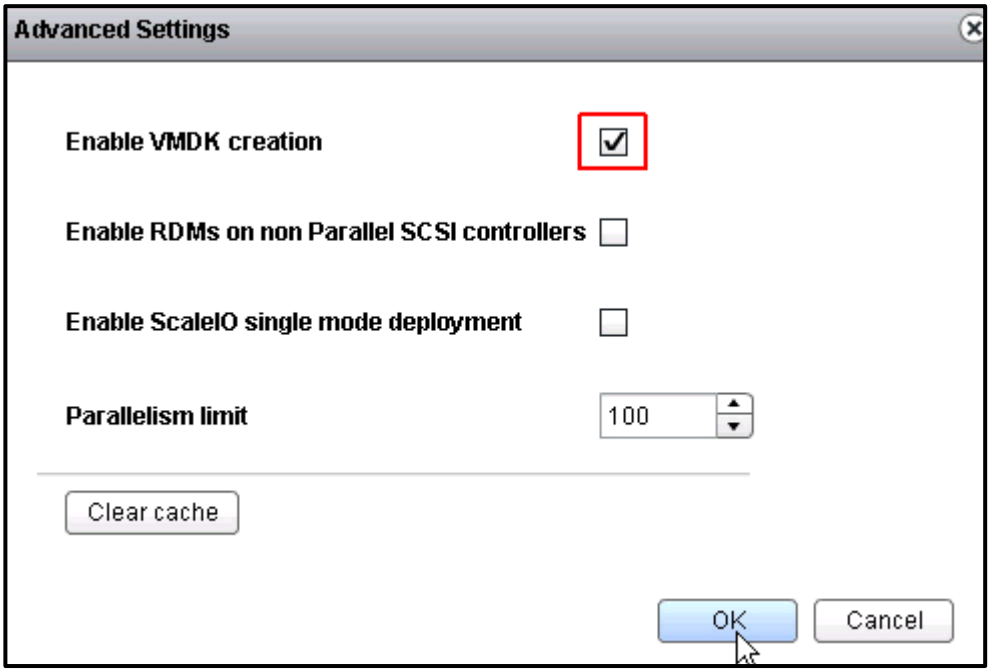 <br> |
| 4. | Minimize the browser, and open a connection to the vCenter appliance directly, by clicking on the icon on the desktop (click Yes if you encounter a security prompt).<br> <br><br>*(it will launch a putty connection and logon as root \| P@ssw0rd, automatically)* |
| 5. | Enter the following command to edit the settings file for the deployment wizard<br>**`vi /etc/vmware/vsphere-client/vc-packages/scaleio/serverConfig.cfg`**<br>`Using username "root".`<br>`VMware vCenter Server Appliance`<br>`Last login: Wed Jul 29 18:42:13 2015 from dc.emc.edu`<br>`vc:~ # vi /etc/vmware/vsphere-client/vc-packages/scaleio/serverConfig.cfg` |

| Step | Action |
|---|---|
| 6. | Insert the following line into the file as shown<br><br>**<vmdk_eager_zeroed>false</vmdk_eager_zeroed>**<br><br>BEFORE<br>`<?xml version="1.0" encoding="UTF-8" standalone="no"?>`<br>`<server_configuration>`<br>`  <general>`<br>`  <allow_vmdk_creation>true</allow_vmdk_creation>`<br>`  <allow_rdm_on_non_parallel_scsi>false</allow_rdm`<br>`  <allow_non_clustered_ScaleIO_system>false</allow`<br>`  <datastore_spare_capacity_percentage>10</datasto`<br>`  <svm_reserved_size_in_gb>8</svm_reserved_size_in`<br>`  <paralellism_limit>100</paralellism_limit>`<br>`  <support_unsigned_vib>false</support_unsigned_vi`<br>`  <vmware_timeout_in_minutes>2</vmware_timeout_in_`<br>`  </general>`<br>`</server_configuration>`<br><br>AFTER<br>`<?xml version="1.0" encoding="UTF-8" standalone="no"?>`<br>`<server_configuration>`<br>`  <general>`<br>`  <vmdk_eager_zeroed>false</vmdk_eager_zeroed>`<br>`  <allow_vmdk_creation>true</allow_vmdk_creation>`<br>`  <allow_rdm_on_non_parallel_scsi>false</allow_rdm`<br>`  <allow_non_clustered_ScaleIO_system>false</allow`<br>`  <datastore_spare_capacity_percentage>10</datasto`<br>`  <svm_reserved_size_in_gb>8</svm_reserved_size_in`<br>`  <paralellism_limit>100</paralellism_limit>`<br>`  <support_unsigned_vib>false</support_unsigned_vi`<br>`  <vmware_timeout_in_minutes>2</vmware_timeout_in_`<br>`  </general>`<br>`</server_configuration>` |
| 7. | Save and exit the file (<Esc> :  w q <Enter> ) |

| Step | Action |
|---|---|
| 8. | Return the ScaleIO Plugin session<br><br>**EMC ScaleIO v1.32-403.2**<br>EMC ScaleIO is software that creates a virtual SAN from local server-based storage to deliver elastic and scalable performance and capacity on demand. ScaleIO converges storage and compute resources, scaling from tens to<br><br>**Basic tasks**<br>Install SDC on ESX<br>Deploy ScaleIO environment<br>Register ScaleIO system |

## Lab 2-B: Task 5 – Deploy ScaleIO onto VMware

| Step | Action |
|---|---|
| | This task has you entering all the values and completing the deployment. Here is a summary:<br><br>- New ScaleIO Cluster with a name of "SIO-VMW" and P@ssw0rd as the admin password<br>  ○ One protection domain "PD1"<br>  ○ One Storage Pool "SP1"<br>  ○ A Gateway VM on 192.168.<br>- MDM<br>  ○ 192.168. → Primary<br>  ○ 192.168.1.31 → Secondary<br>  ○ 192.168.1.32 → Tie Breaker<br>- SDS<br>  ○ All three ESXi servers<br>  ○ 1 disk added from each SDS (105GB, VMDK created)<br>- SDC<br>  ○ All three ESXi servers<br>  ○ SCSI LUN number comparison Disabled<br>- Networking<br>  ○ ScaleIO management traffic on existing 192.168.1.x network<br>  ○ Data traffic on new 192.168.2.x network |

| Step | Action |
|---|---|
| 1. | From the ScaleIO Plugin page, click Deploy ScaleIO environment<br> <br> |
| 2. | On the Select Installation Type screen, leave at the default (Create new ScaleIO system), and click Next<br> 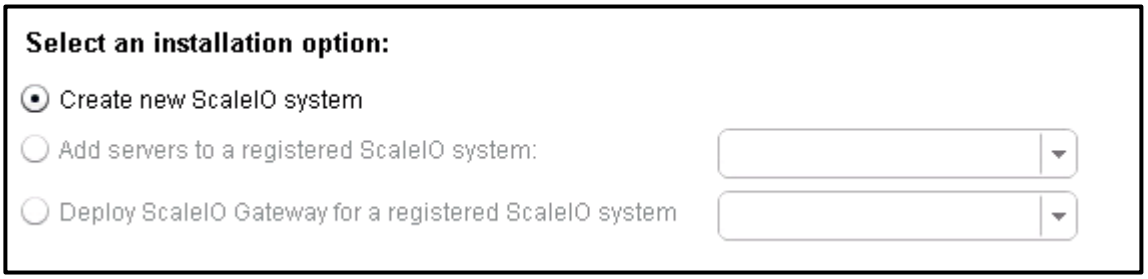 <br> |
| 3. | On the Confirm license screen, Click I agree… and click Next<br> 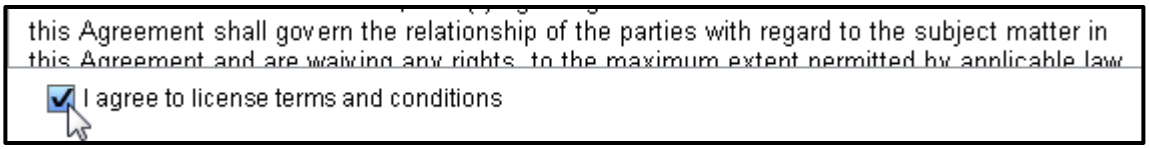 <br> |
| 4. | On the Create new system screen, enter SIO-VMW as the system name, enter and confirm P@ssw0rd as the password, and click Next<br> 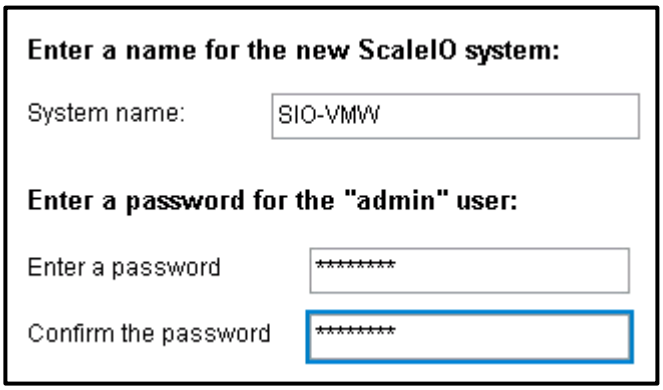 <br> |

| Step | Action |
|---|---|
| 5. | On the Add ESX hosts to cluster screen, Select vCenter as the vCenter server, select all three ESX hosts (you can select any parent object to do this as well), and click Next.<br> 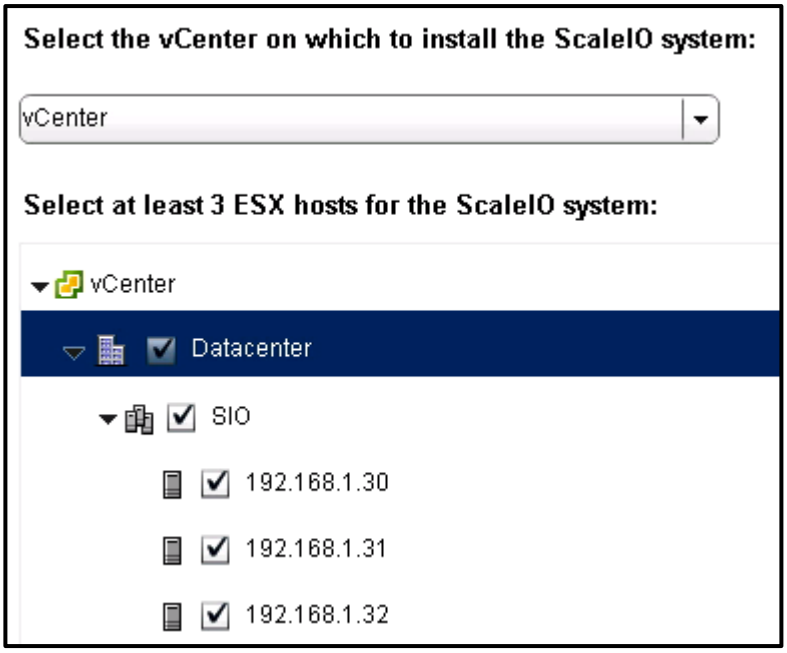 <br> |
| 6. | On the Select management components screen, assign the roles as shown below, and click Next<br> 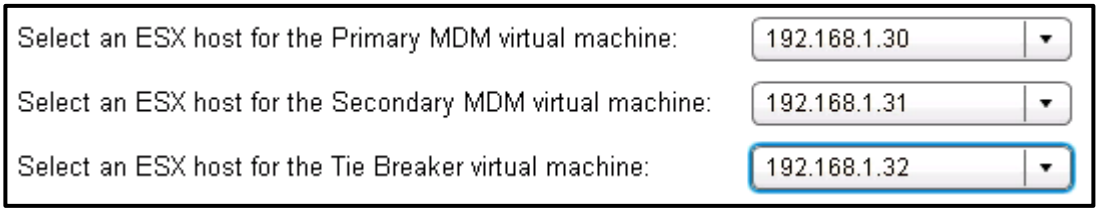 <br> |
| 7. | On the Configure Call Home screen, remove the check from Configure Call Home, and click Next<br> 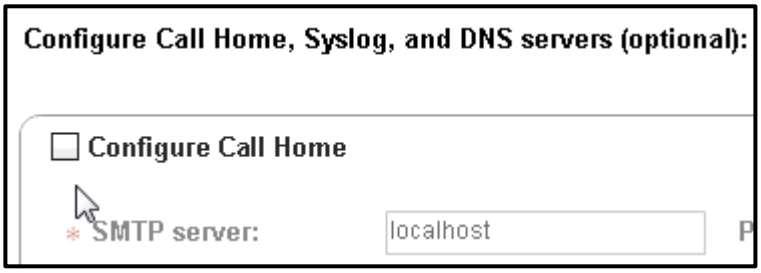 <br> |

| Step | Action |
|---|---|
| 8. | On the Configure protection domain screen, Enter PD1, click Add (it will appear below), and click Next.<br> 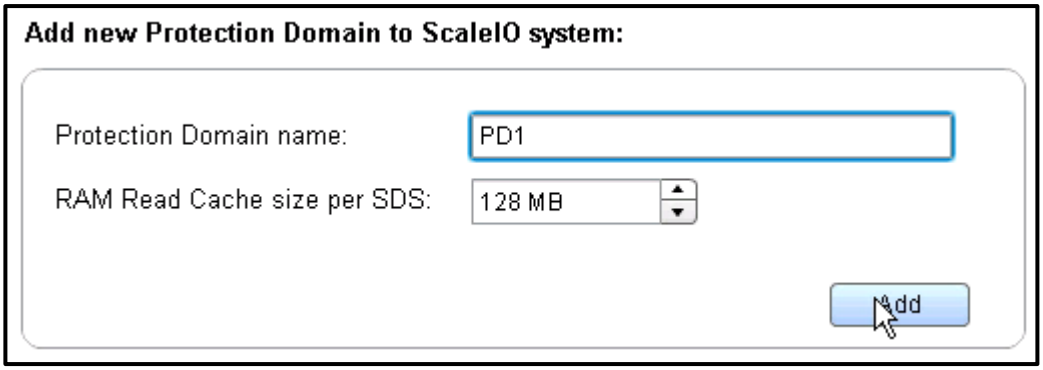 <br><br> 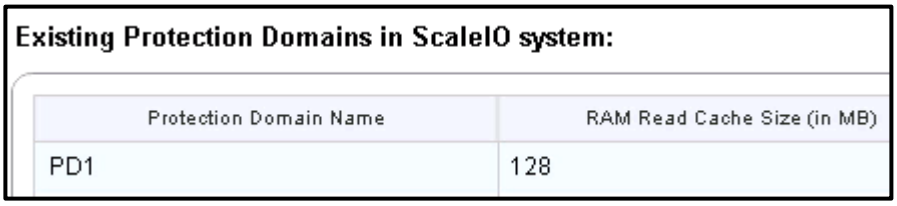 <br> |
| 9. | On the Configure Storage Pools screen, enter SP1 as the storage pool name (under protection domain PD1), click Add (it will appear below), and click Next.<br> 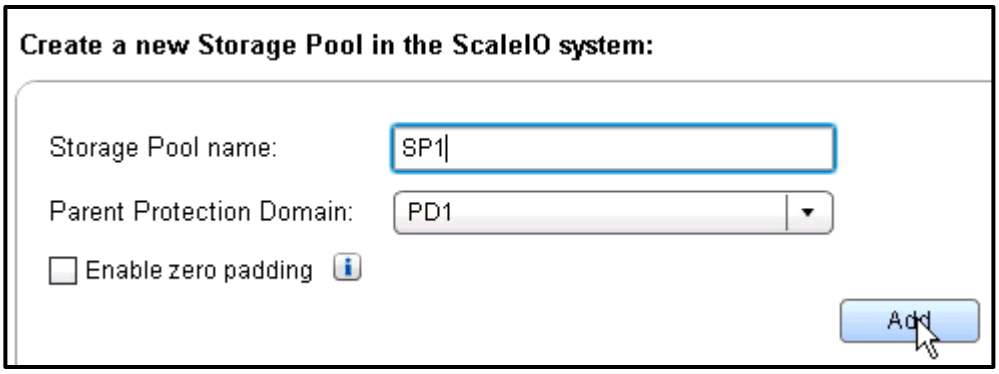 <br><br> 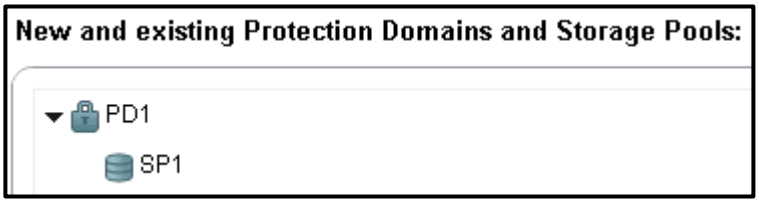 <br> |

| Step | Action |
|---|---|
| 10. | On the Create Fault Sets screen, do not enter a fault set, and click Next<br> 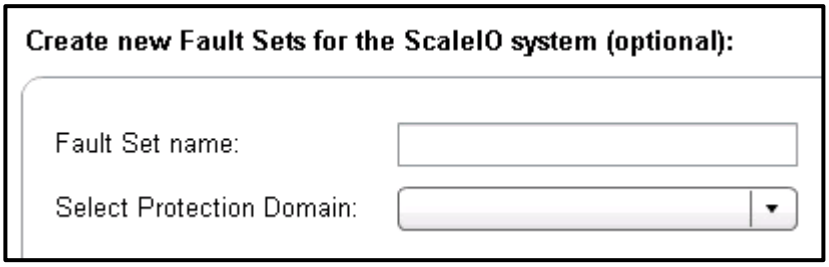 <br> |
| 11. | On the Add SDS screen, select all three ESX hosts (can also select a parent object), and click Next.<br> 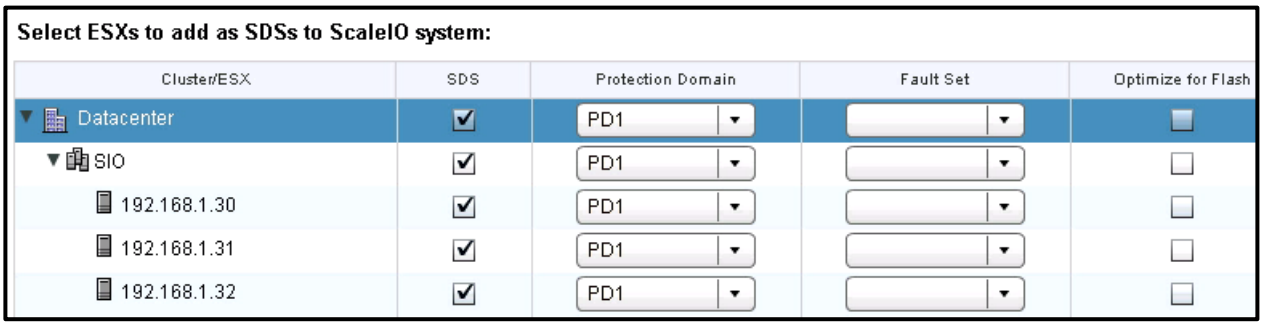 <br> |
| 12. | On the Add devices to SDS screen, select the first ESX host (192.168. ) as shown. Do NOT click Next.<br> 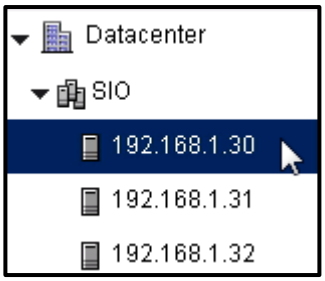 <br> |

| Step | Action |
|---|---|
| 13. | Click on the Select Devices tab, select the C0:T1:L0 device, click Add Device and Create VM… (VMDK). Do not click Next. |

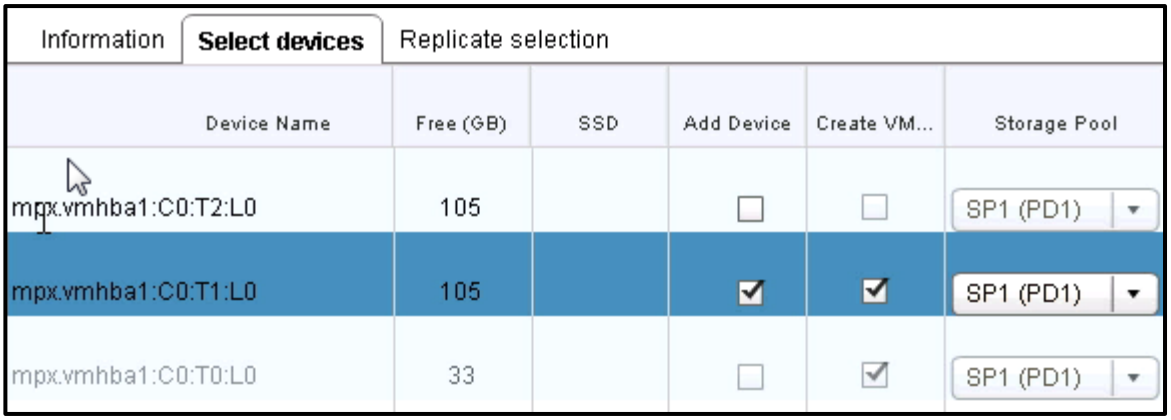

| Step | Action |
|---|---|
| 14. | Repeat the previous two steps for 192.168.1.31, and 192.168.1.32. Then, click Next. |

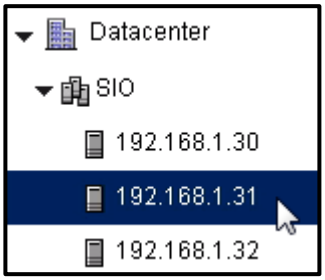


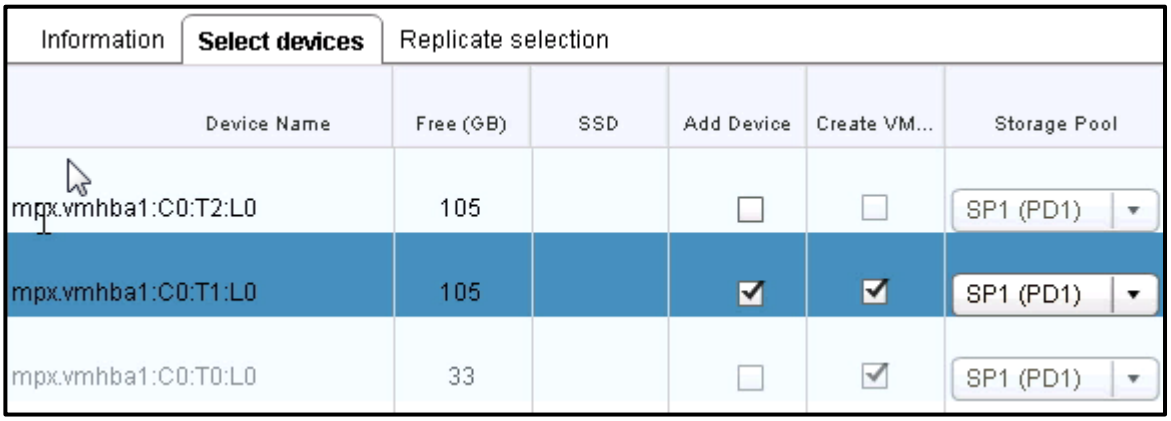


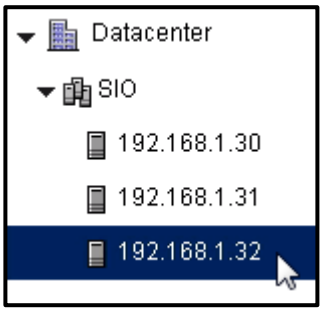


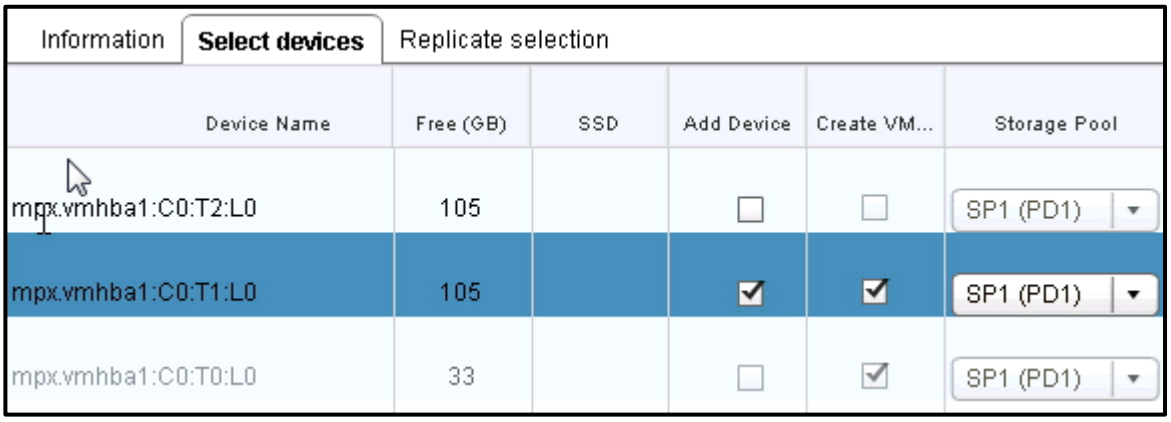

| Step | Action |
|---|---|
| 15. | On the Add SDC screen, select all three hosts (or select a parent object), enter P@ssw0rd as the password for all (click Show passwords to make sure it is typed correctly). Do NOT click Next.<br> 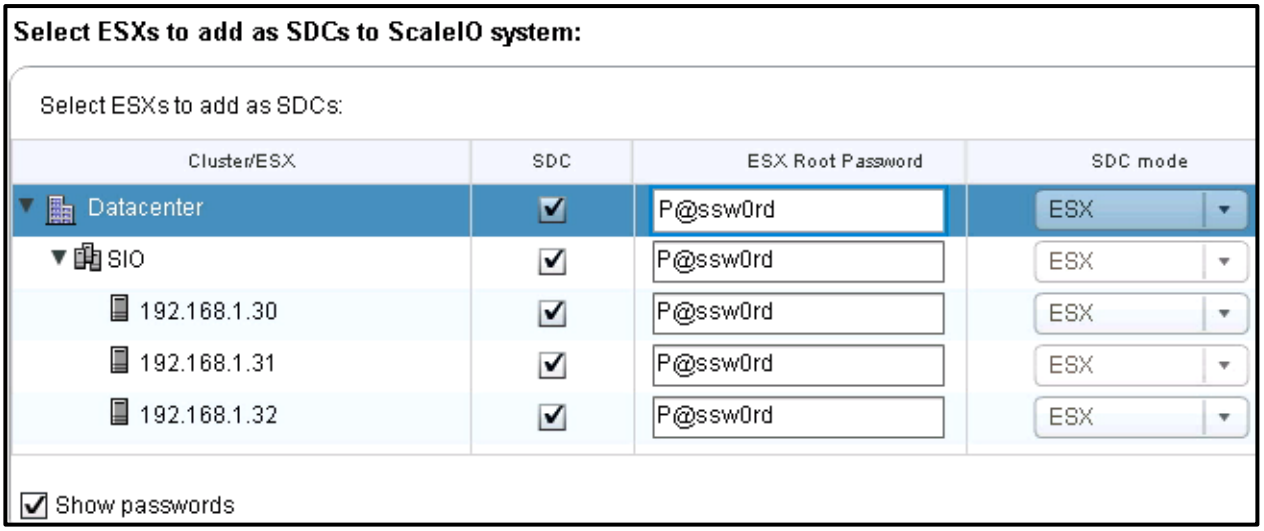 <br> |
| 16. | On the same screen, click Disable for Enable/Disable SCSI LUN number comparison… Click OK on the popup, then click Next.<br> 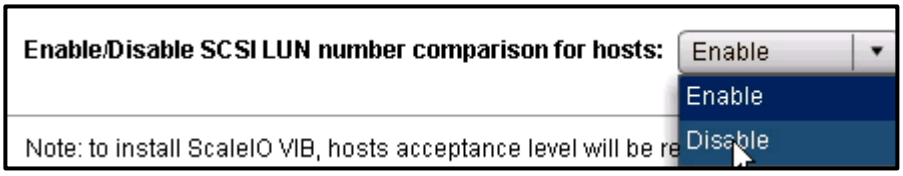 <br><br> 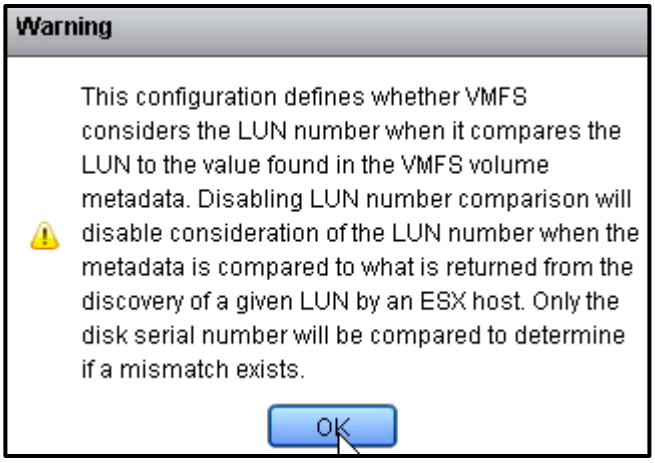 <br> |

| Step | Action |
|---|---|
| 17. | On the Configure Upgrade Components screen, select 192.168.  as the host for the Gateway VM, enter P@ssw0rd as the password for the gateway, enter P@ssw0rd as the password for LIA, then click Next.<br> 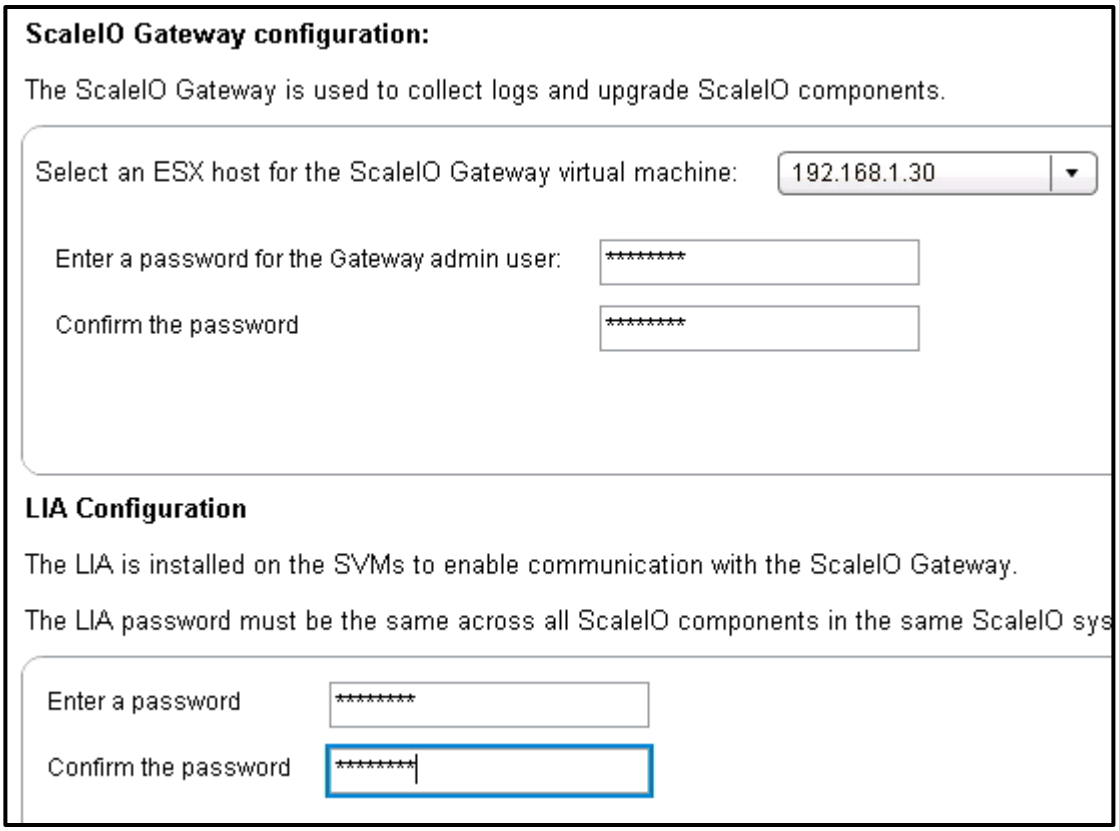 <br> |

| Step | Action |
|---|---|
| 18. | On the Select OVA template screen, choose the previously uploaded template, enter P@ssw0rd as the root password, and click Next<br> 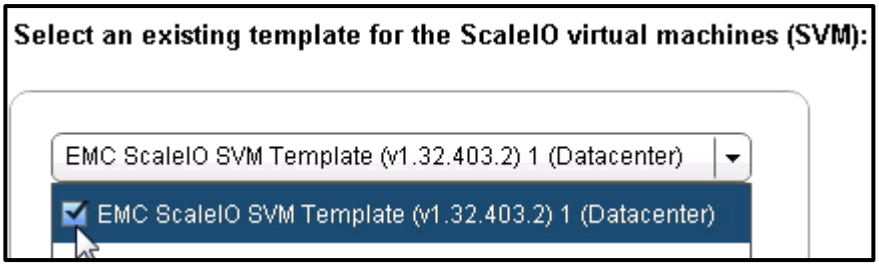 <br><br> 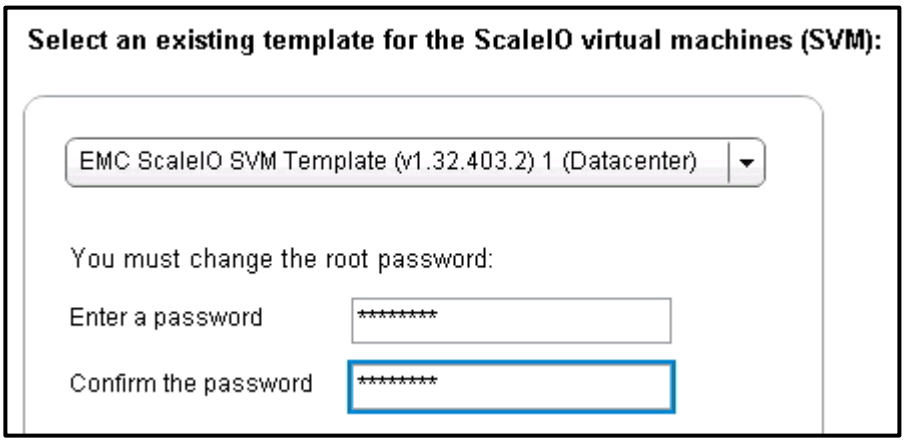 <br> |
| 19. | On the Configure networks screen, select the VM Network for the Management network. For the Data network, click Create new network<br> 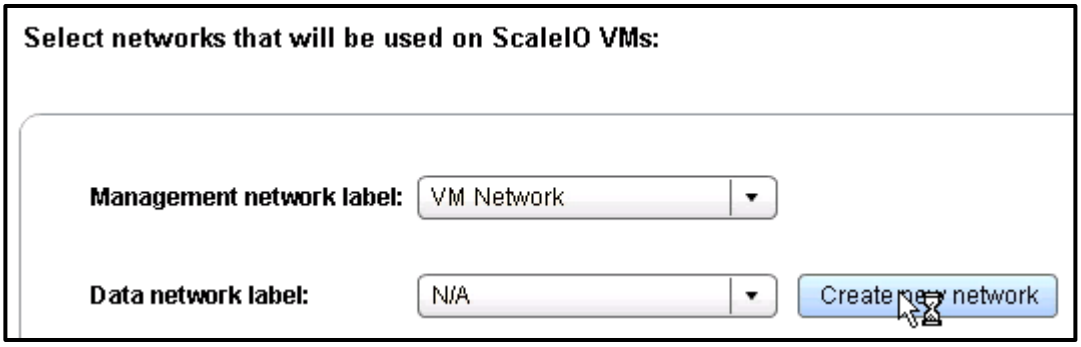 <br> |

| Step | Action |
|---|---|
| 20. | For the values for the new network, enter what is shown below, then click OK<br> 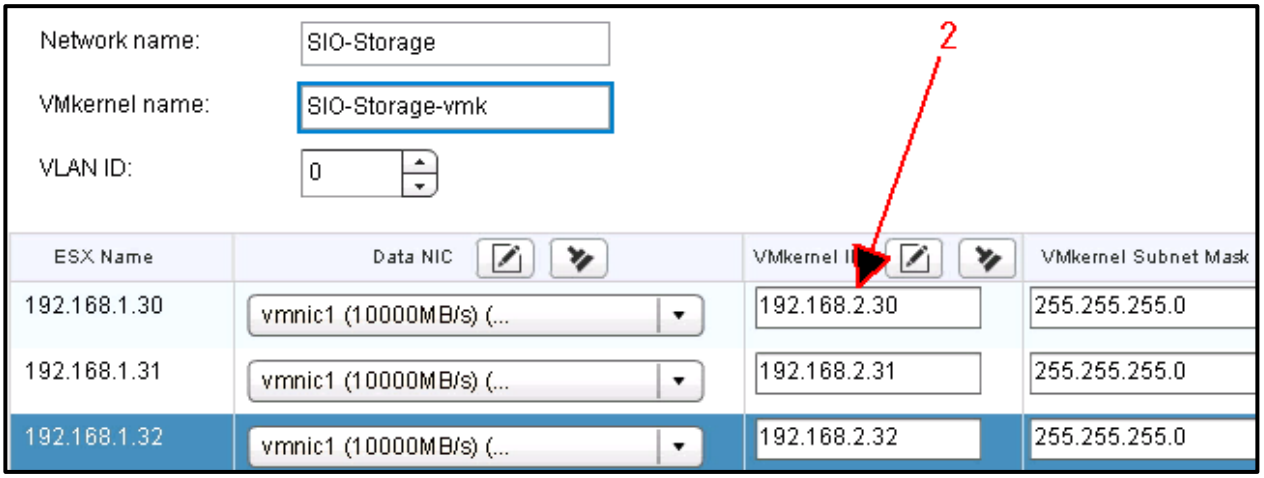 <br> |
| 21. | The network configuration will take place.<br><br>When it is finished configuring the networks, click Close on the progress popup, and close on the Create New Data Network window, to return to the Configure networks screen.<br> 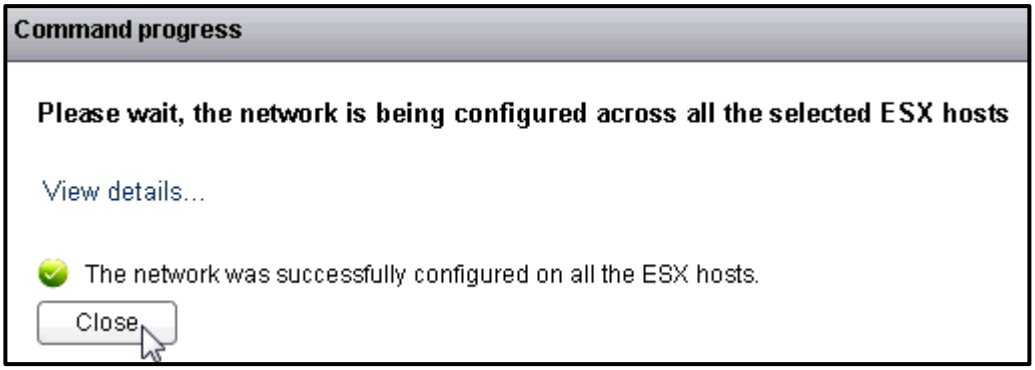 <br><br> <br> |

| Step | Action |
|---|---|
| 22. | Wait a moment for the new SIO-Storage network to appear as the Data network. Then click Next.<br> 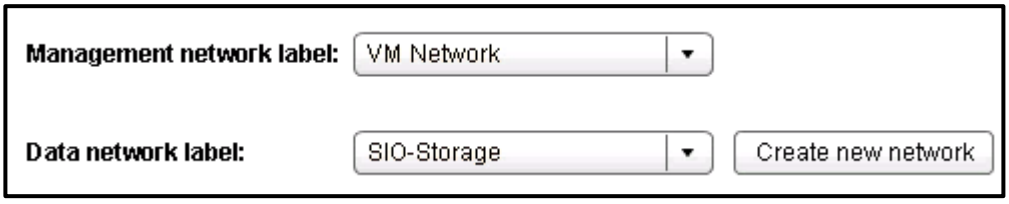 <br> |
| 23. | On the Configure SVM Screen, enter the values as shown below (scroll from left to right to access the different values), and then click Next.<br> 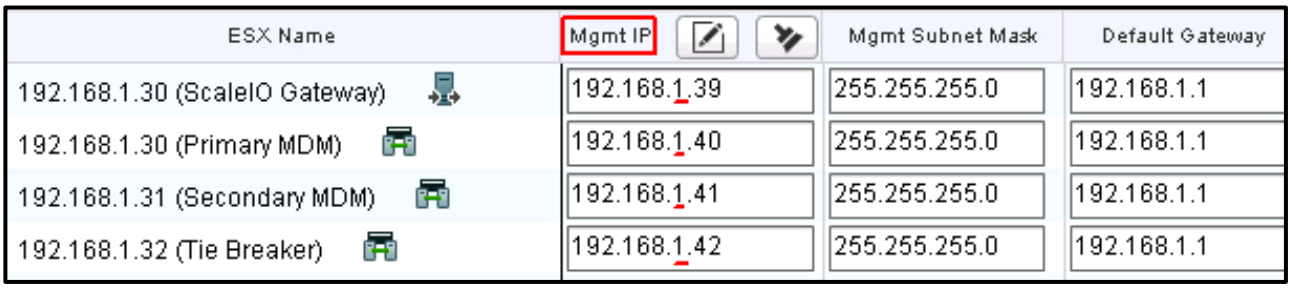 <br><br> 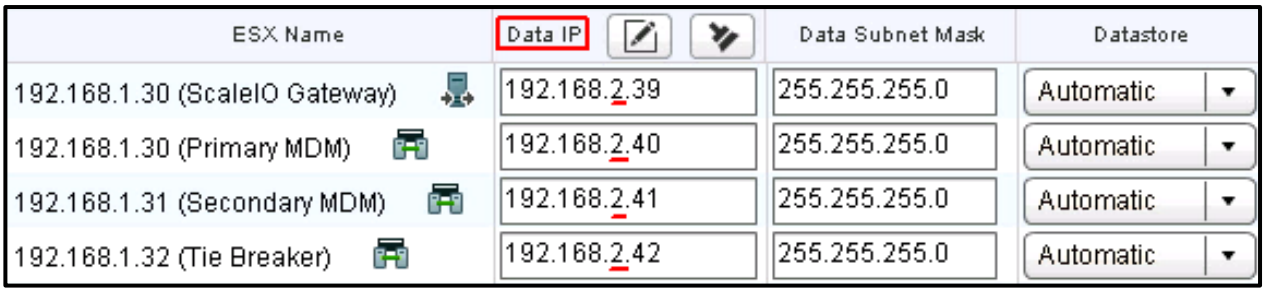 <br> |

| Step | Action |
|---|---|
| 24. | Review the configuration summary, click Finish.<br><br>**Please review configuration summary before installation begins:**<br><br>Number of SDSs: 3<br>Number of SDCs: 3<br>Number of ESXs: 3<br>Number of physical devices: 3<br>Number of Protection Domains: 1<br>Number of Storage Pools: 1<br>Total capacity (GB): 315.00<br><br>Click Finish to begin deployment or Back to make changes. |
| 25. | On the authorization popup, enter root \| P@ssw0rd, and click OK.<br><br>**vCenter Authorization**<br><br>vCenter authentication required<br><br>Enter user name and password to connect to vCenter vCenter:<br><br>User name: root<br>Password: ********<br><br>OK Cancel |

| Step | Action |
|---|---|
| 26. | The installation will begin. It may take up to 45 minutes to complete. It is common in our environment to experience errors. If you see an Error, click Pause (then Yes to confirm), wait for it to stop the Deployment, click Continue the deployment (Yes to confirm)<br> <br><br> <br><br> 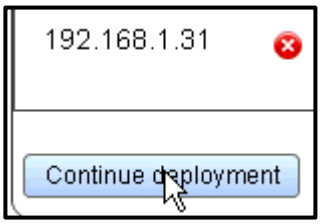 <br> |
| 27. | Do not be concerned with errors regarding Datastore usage. This is due to “VMDK Creation” which creates a single VMDK on each disk, consuming nearly the whole disk.<br> 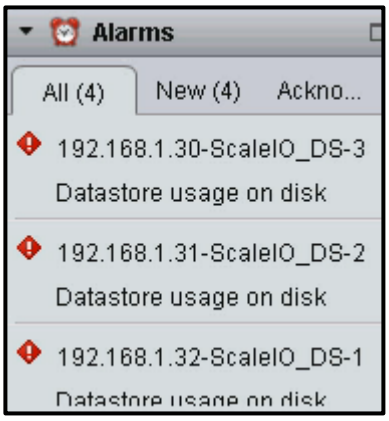 <br> |

| Step | Action |
|---|---|
| 28. | Also, these errors are expected and can be ignored<br> <br> |
| 29. | When it finally completes without errors, click Finish<br> 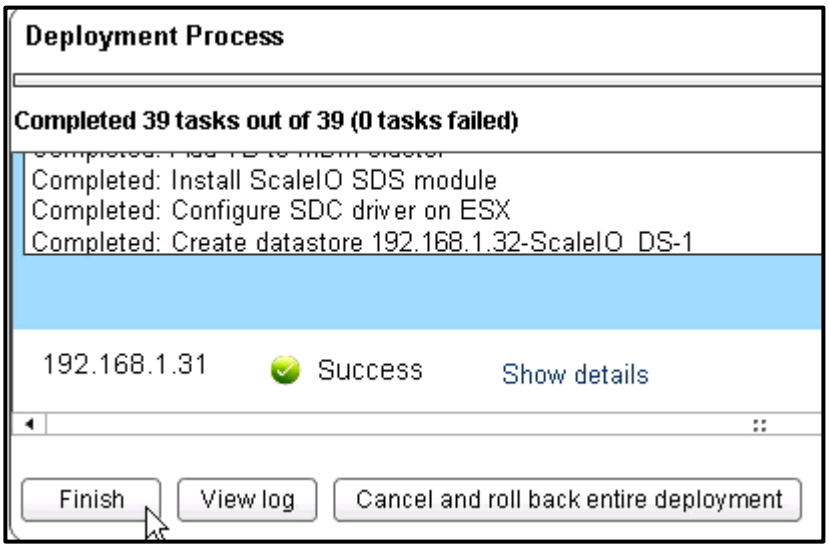 <br> |

| Step | Action |
|---|---|
| 30. | The lower-left of the vSphere screen should look like this after a successful deployment.<br> 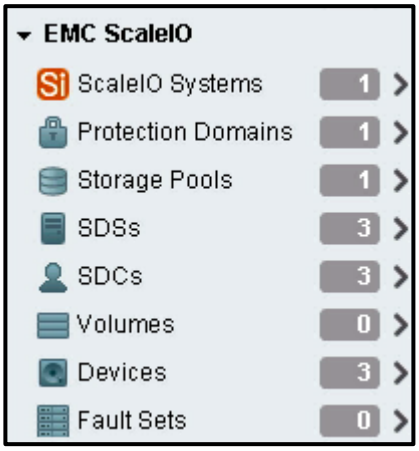 <br> |

## Lab Exercise 3-A: ScaleIO Configuration and Management

| | |
|---|---|
| **Purpose:** | Perform basic configuration, management and monitoring of a ScaleIO cluster in a physical Linux/Windows environment. |

| | |
|---|---|
| **Tasks:** | ✓ Practice using the ScaleIO GUI to perform the supported configuration, management and monitoring operations<br>✓ Gain familiarity with the ScaleIO CLI syntax by using it to exercise typical day-to-day tasks such as provisioning and storage management |

| | |
|---|---|
| **References:** | ScaleIO Lab Diagram (graphically showing relevant details for all VMs within a student pod – refer to Page 2) |

## Lab 3-A: Task 1 – GUI Introduction

| Step | Action |
|---|---|
| | This task has you taking a quick look at the ScaleIO GUI. It is available for Linux and Windows. You will install it on the Windows DC, and take a quick look at the dashboard and backend views. |
| 1. | From the "dc" workstation, open an explorer window to c:\share\sw\GUI_for_Windows.<br><br>Launch the msi installer<br><br> 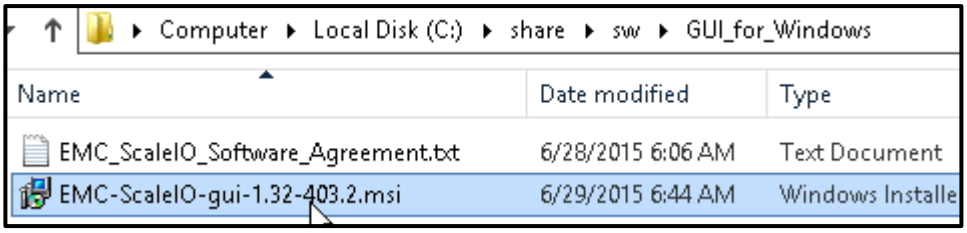 |

| Step | Action |
|---|---|
| 2. | Accept all the default settings and Finish<br> |
| 3. | Launch the GUI under Start → Programs → ScaleIO<br> 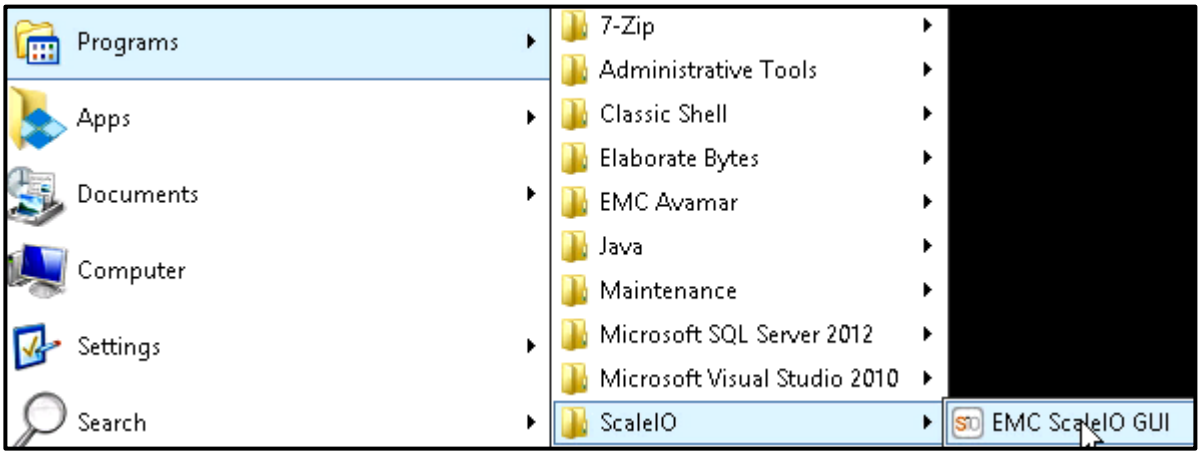 |
| 4. | Log on to lnx-1 as User: admin \| Password: P@ssw0rd<br> 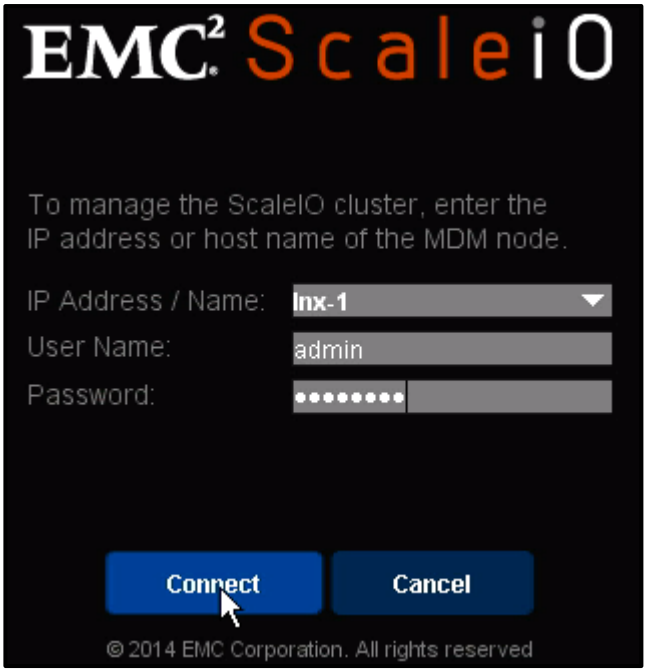 |

| Step | Action |
|---|---|
| 5. | Take a close look at the Dashboard<br> 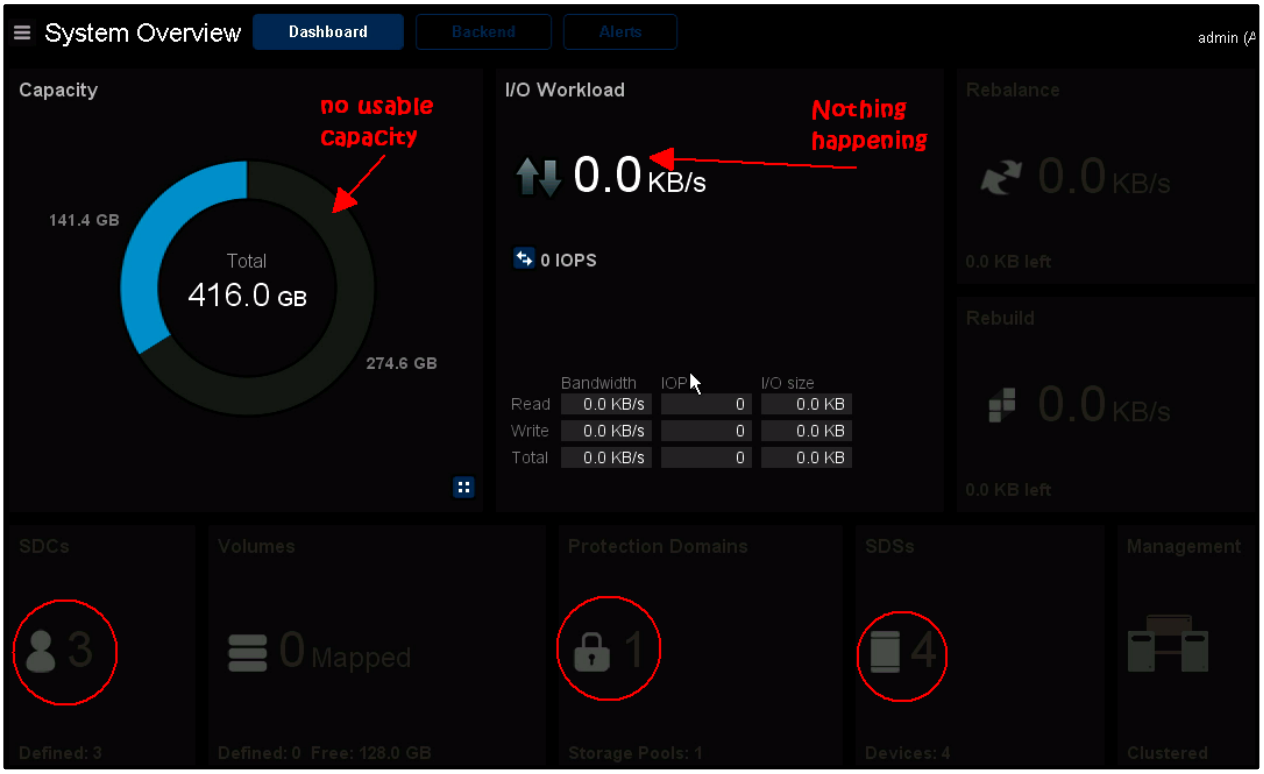 <br> |
| 6. | Click on Backend to take a close look at the backend view<br> 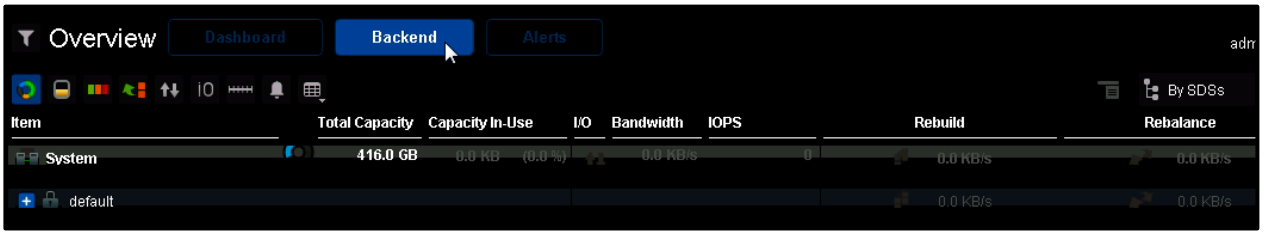 <br> |

| Step | Action |
|---|---|
| 7. | The default view is By SDS. Take a look at the SDS's in the protection domain by expanding the "default" protection domain<br> 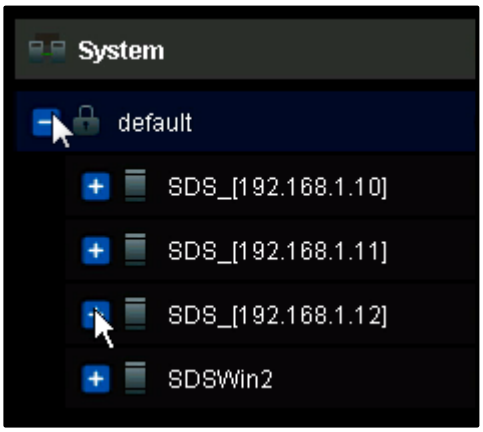 <br> |
| 8. | In the upper-right of the interface, switch from By SDS, to By Storage Pools.<br> <br> |
| 9. | Expand the tree on left to view all the components comprising the current default storage pool<br> 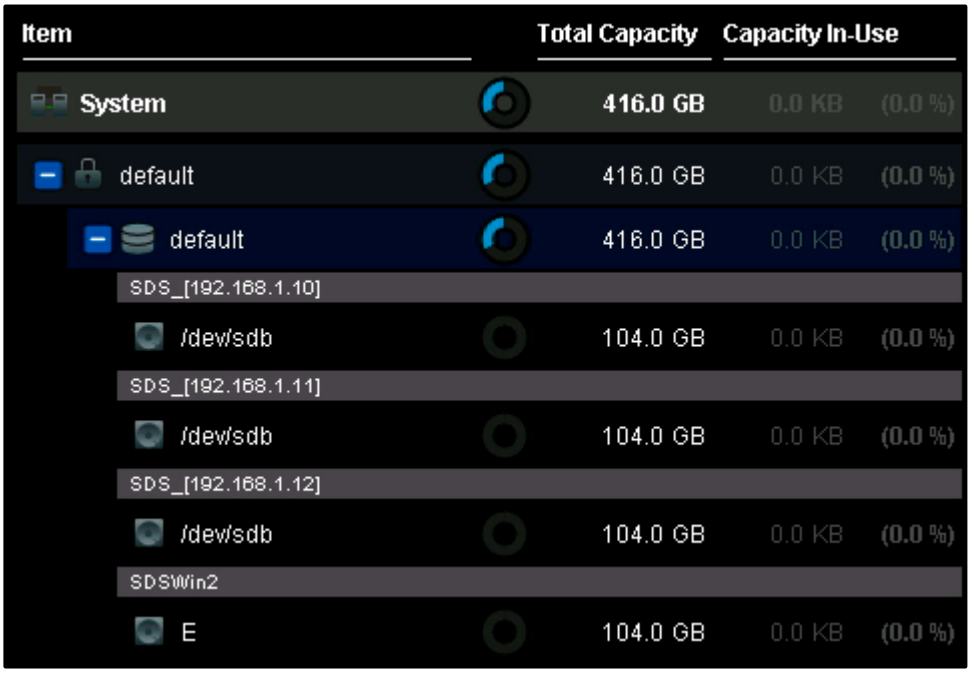 <br> |

## Lab 3-A: Task 2 – SCLI

| Step | Action |
|---|---|
| | Here you will run scli commands to gather basic information. |
| 1. | Minimize the GUI, and launch a putty session with lnx-1, by clicking on the desktop icon (will log you in as root)<br> <br><br>`Using username "root".`<br>`Last login: Thu Jul 30 10:36:45 2015 from 192.168.1.50`<br>`[root@Lnx-1 ~]#` |
| 2. | Run the following command to display basic help information on<br>**`scli`**<br>`[root@Lnx-1 ~]# scli`<br>`Usage: scli --help [Options]`<br>`Description: Display help about one or all types of commands`<br>`Parameters:`<br>`Options: CHOOSE ONE`<br>`--mdm  Display help about commands relating to the MDM`<br>`--sds  Display help about commands relating to the SDS`<br>`--vol  Display help about commands relating to volumes`<br>`--sdc  Display help about commands relating to SDC`<br>`--SCSI  Display help about commands relating to SCSI`<br>`--general  Display help about general commands, e.g. licen`<br>`on, etc.`<br>`--user  Users login and management`<br>`--all  Display help about all commands` |
| 3. | Run the following command to gain privileged access to the ScaleIO cluster<br>**`scli --login --username admin --password P@ssw0rd`**<br>`[root@Lnx-1 ~]# scli --login --username admin --password P@ssw0rd`<br>`Logged in. User role is Administrator. System ID is 4b36f8677d4ecb5d`<br>`[root@Lnx-1 ~]#`<br>*Note: if you omit the "--password" switch, it will prompt for secure password entry.* |

| Step | Action |
|---|---|
| 4. | Enter the following command to to retrieve lots of information from the cluster<br><br>**scli –query_all**<br><br>**[root@Lnx-1 ~]# scli --query_all**<br><br>The output will quickly scroll past the screen |
| 5. | To retrieve the information one screen at a time, enter the following command:<br><br>**scli –query_all \| less**<br><br>**[root@Lnx-1 ~]# scli --query_all \| less** |
| 6. | It will output information just one screen at a time. Hit [spacebar] to continue to the next page and q to exit the command.<br><br>System Info:<br>Product: EMC ScaleIO Version: R1_32.403.2<br>ID: 4b36f8677d4ecb5d<br>Manager ID: 0000000000000000<br><br>License info:<br>Installation ID: 370409a270d6e920<br>SWID:<br>Maximum capacity: Unlimited<br>Usage time left: Unlimited *** Non-Production License ***<br>Enterprise features: Enabled<br>The system was activated 1 days ago |
| 7. | Enter the following command to identify the protection domain.<br><br>**scli --query_all \| grep Protection**<br><br>[root@Lnx-1 ~]# scli --query_all \| grep Protection<br>Query all returned 1 Protection Domain:<br>Protection Domain default (Id: 18cafc9300000000) has 1 storage pools, 0 Fault Sets, 4 SDS nodes, 0 v<br>olumes and 128.0 GB (131072 MB) available for volume allocation |

| Step | Action |
|---|---|
| 8. | Enter the following command for detailed information on the protection domain "default". (You can either scroll through the results or add "`\| less`" to the command to view one page at a time. Take a moment to review the output.<br><br>**`scli --query_protection_domain --protection_domain_name default`**<br><br>`[root@Lnx-1 ~]# scli --query_protection_domain --protection_domain_name default` |
| 9. | Enter the following command to identify the protection domain.<br><br>**`scli --query_all \| grep Storage`**<br><br>`[root@Lnx-1 ~]# scli --query_all \| grep Storage`<br>`Storage Pool default (Id: a6041e4d00000000) has 0 volumes and 128.0 GB (131072 MB)`<br>`ume allocation`<br>`[root@Lnx-1 ~]#` |
| 10. | Enter the following command (on a single line) for detailed information on the Storage Pool "default". (You can either scroll through the results or add "`\| less`" to the command to view one page at a time. Take a moment to review the output.<br><br>**`scli --query_storage_pool --protection_domain_name default --storage_pool_name default`**<br><br>`[root@Lnx-1 ~]# scli --query_storage_pool --protection_domain_name default --storage_pool_name defau`<br>`lt` |
| 11. | Run the following command to list your SDSs:<br><br>**`scli --query_all_sds`**<br><br>`Query-all-SDS returned 4 SDS nodes.`<br><br>`Protection Domain 18cafc9300000000 Name: default`<br>`SDS ID: b8d54f6d00000004 Name: SDSWin2 State: Connected, Joined IP: 192.168.1.21 Port: 7072`<br>`SDS ID: b8d5285f00000002 Name: SDS_[192.168.1.11] State: Connected, Joined IP: 192.168.1.11 Port: 70`<br>`72`<br>`SDS ID: b8d5285e00000001 Name: SDS_[192.168.1.12] State: Connected, Joined IP: 192.168.1.12 Port: 70`<br>`72`<br>`SDS ID: b8d5285d00000000 Name: SDS_[192.168.1.10] State: Connected, Joined IP: 192.168.1.10 Port: 70`<br>`72`<br>`[root@Lnx-1 ~]#` |

| Step | Action |
|---|---|
| 12. | Query one of your SDSs, using its ID and carefully examine the output:<br><br>**scli --query_sds --sds_id <SDS ID>**<br><br>SDS ID: b8d54f6d00000004 Name: SDSWin2 State: Connected, Jo<br>SDS ID: b8d5285f00000002 Name: SDS_[192.168.1.11] State: Co<br>72<br>SDS ID: b8d5285e00000001 Name: SDS_[192.168.1.12] State: Co<br>72<br>SDS ID: b8d5285d00000000 Name: SDS_[192.168.1.10] State: Co<br>72<br>[root@Lnx-1 ~]# scli --query_sds --sds_id b8d54f6d00000004<br><br>Look for information like, Path: and total and unused capacity. |
| 13. | Run this command to list your SDCs:<br><br>**scli --query_all_sdc**<br><br>Example output:<br><br>[root@Lnx-1 ~]# scli --query_all_sdc<br>MDM restricted SDC mode: Disabled<br>Query all SDC returned 3 SDC nodes.<br>SDC ID: c842438d00000000 Name: N/A IP: 192.168.1.10<br>EACBA3939F1<br>Read bandwidth: 0 IOPS 0 Bytes per-second<br>Write bandwidth: 0 IOPS 0 Bytes per-second<br>SDC ID: c842438e00000001 Name: N/A IP: 192.168.1.12<br>D6A4CB38FAF<br>Read bandwidth: 0 IOPS 0 Bytes per-second<br>Write bandwidth: 0 IOPS 0 Bytes per-second<br>SDC ID: c842438f00000002 Name: N/A IP: 192.168.1.11<br>EDE4BA577F3<br>Read bandwidth: 0 IOPS 0 Bytes per-second<br>Write bandwidth: 0 IOPS 0 Bytes per-second |

| Step | Action |
|---|---|
| 14. | All three of the SDCs still have the name given to them by the Installation Manager during their creation. Name them to reflect their purpose and host names.<br><br>**`scli --rename_sdc --sdc_id <SDC ID> --new_name SDC_Lnx1`**<br><br>(pick the SDC ID corresponding to the IP 192.168.1.10)<br><br>**`scli --rename_sdc --sdc_id <SDC ID> --new_name SDC_Lnx2`**<br><br>(pick the SDC ID corresponding to the IP 192.168.1.11)<br><br>**`scli --rename_sdc --sdc_id <SDC ID> --new_name SDC_Lnx3`**<br><br>(pick the SDC ID corresponding to the IP 192.168.1.12)<br><br>`[root@Lnx-1 ~]# scli --rename_sdc --sdc_id c842438d00000000 --new_name SDC_Lnx1`<br>`Successfully renamed sdc with ID c842438d00000000 to SDC_Lnx1`<br>`[root@Lnx-1 ~]# scli --rename_sdc --sdc_id c842438d00000002 --new_name SDC_Lnx2`<br>`Error: MDM failed command.  Status: Could not find the SDC`<br>`[root@Lnx-1 ~]# scli --rename_sdc --sdc_id c842438f00000002 --new_name SDC_Lnx2`<br>`Successfully renamed sdc with ID c842438f00000002 to SDC_Lnx2`<br>`[root@Lnx-1 ~]# scli --rename_sdc --sdc_id c842438e00000001 --new_name SDC_Lnx3`<br>`Successfully renamed sdc with ID c842438e00000001 to SDC_Lnx3` |
| 15. | Rerun **`scli --query_all_sdc`** and ensure the new names are in effect:<br><br>`[root@Lnx-1 ~]# scli --query_all_sdc`<br>`MDM restricted SDC mode: Disabled`<br>`Query all SDC returned 3 SDC nodes.`<br>`SDC ID: c842438d00000000 Name: SDC_Lnx1 IP: 192.168.1.10`<br>`214-0EACBA3939F1`<br>`  Read bandwidth:  0 IOPS 0 Bytes per-second`<br>`  Write bandwidth:  0 IOPS 0 Bytes per-second`<br>`SDC ID: c842438e00000001 Name: SDC_Lnx3 IP: 192.168.1.12`<br>`B0E-5D6A4CB38FAF`<br>`  Read bandwidth:  0 IOPS 0 Bytes per-second`<br>`  Write bandwidth:  0 IOPS 0 Bytes per-second`<br>`SDC ID: c842438f00000002 Name: SDC_Lnx2 IP: 192.168.1.11`<br>`94C-6EDE4BA577F3`<br>`  Read bandwidth:  0 IOPS 0 Bytes per-second`<br>`  Write bandwidth:  0 IOPS 0 Bytes per-second` |

## Lab 3-A: Task 3 – Create volumes

| Step | Action |
|---|---|

| Step | Action |
|---|---|
| | You will create volumes via scli. |
| 1. | At the command prompt on Lnx-1, enter (all one command):<br><br>**scli --add_volume --protection_domain_name default --storage_pool_name default --thin_provisioned --volume_name vol1 --size_gb 12**<br><br>[root@Lnx-1 ~]# scli --add_volume --protection_domain_name default --storage_pool_name default --thin_provisioned --volume_name vol1 --size_gb 12<br>Warning: Rounding up the volume size to 16 GB<br>Successfully created volume of size 16 GB. Object ID b8fb1c3900000000<br>[root@Lnx-1 ~]#<br><br>**Note:** ScaleIO rounded up the provisioned capacity to 16 GB. Due to the nature of ScaleIO' s mesh mirroring, it works best in units of 8 GB. |
| 2. | List all the volumes and look for the volume you just created.<br><br>**scli --query_all_volumes**<br><br>[root@Lnx-1 ~]# scli --query_all_volumes<br>Query-all-volumes returned 1 volumes<br>Protection Domain 18cafc9300000000 Name: default<br>Storage Pool a6041e4d00000000 Name: default<br>Volume ID: b8fb1c3900000000 Name: vol1 Size: 16.0 GB (16384 MB) Unmapped Thin-provisioned |
| 3. | Minimize the Putty session and bring up the ScaleIO GUI. |
| 4. | On **Dashboard→System Overview**, look for the **Volumes** tile. You should see 1 volume defined.<br><br> 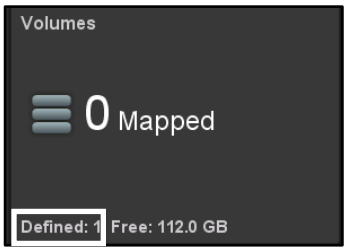 <br> |

| Step | Action |
|---|---|
| 5. | Now look for the **Capacity** tile. Notice, however, that it doesn't show any storage as being used (Protected) yet.<br> 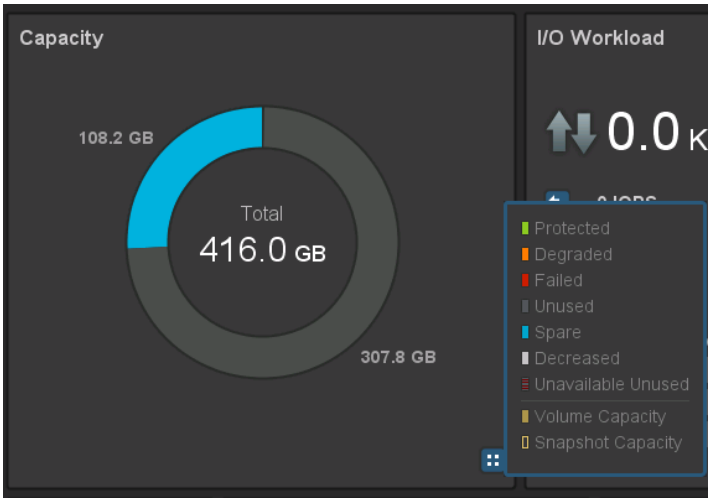 <br><br>Why? _____ |
| 6. | Create a new thin-provisioned volume of 16 GB named vol 2.<br>**`scli --add_volume --protection_domain_name default --storage_pool_name default --thin_provisioned --volume_name vol2 --size_gb 16`**<br>`[root@Lnx-1 ~]# scli --add_volume --protection_domain_name default --storage_pool_name default --thin_provisioned --volume_name vol2 --size_gb 16`<br>Refer back to the GUI. Any change on the **Capacity** circle yet? _____ |
| 7. | Now create two more <u>thick-provisioned</u> volumes, **vol3** and **vol4**.<br>**Note:** Thick provisioning is the default. Notice the lack of the **`--thin_provisioned`** qualifier in the commands.<br>**`scli --add_volume --protection_domain_name default --storage_pool_name default --volume_name vol3 --size_gb 16`**<br>**`scli --add_volume --protection_domain_name default --storage_pool_name default --volume_name vol4 --size_gb 16`**<br>`[root@Lnx-1 ~]# scli --add_volume --protection_domain_name default --storage_pool_name default --volume_name vol3 --size_gb 16`<br>`Successfully created volume of size 16 GB. Object ID b8fb1c3b00000002`<br>`[root@Lnx-1 ~]# scli --add_volume --protection_domain_name default --storage_pool_name default --volume_name vol4 --size_gb 16`<br>`Successfully created volume of size 16 GB. Object ID b8fb1c3c00000003` |

| Step | Action |
|---|---|
| 8. | Check the **Capacity** circle.<br> 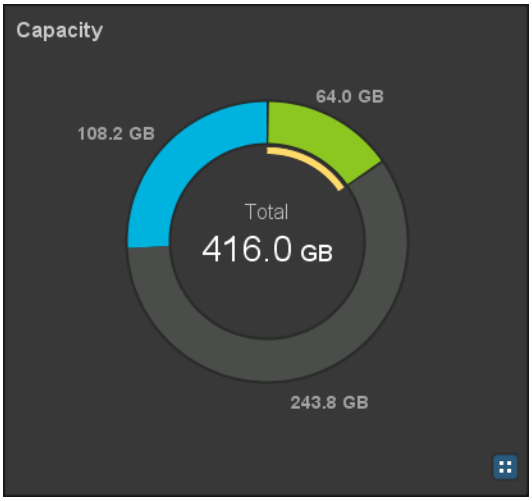 <br><br>You created two 16GB thick-provisioned devices, for a total of 32GB usable capacity.<br><br>Why does ScaleIO show 64GB of protected storage?<br><br>_____ |

## Lab 3-A: Task 4 – Map Volumes to SDC

| Step | Action |
|---|---|
| | You will provision created volumes to an sdc host. Then, from the host, you will run the commands to access the volume. |
| 1. | **Note:** This task assumes that Lnx-1 has ownership of the Primary MDM.<br><br>On Lnx-1, run the following command to map vol1 to Lnx-1:<br><br>**`scli --map_volume_to_sdc --volume_name vol1 --sdc_name SDC_Lnx1`**<br><br>`[root@Lnx-1 ~]# scli --map_volume_to_sdc --volume_name vol1 --sdc_name SDC_Lnx1`<br>`Successfully mapped volume vol1 to SDC SDC_Lnx1`<br>`[root@Lnx-1 ~]#` |

| Step | Action |
|---|---|
| 2. | Still on Lnx-1, (the Primary MDM) map vol2 to Lnx-2 and vol3 to Lnx-3.<br><br>`[root@Lnx-1 ~]# scli --map_volume_to_sdc --volume_name vol2 --sdc_name SDC_Lnx2`<br>`Successfully mapped volume vol2 to SDC SDC_Lnx2`<br>`[root@Lnx-1 ~]# scli --map_volume_to_sdc --volume_name vol3 --sdc_name SDC_Lnx3`<br>`Successfully mapped volume vol3 to SDC SDC_Lnx3`<br>`[root@Lnx-1 ~]#` |
| 3. | **Note:** At this point, we are running commands on Lnx-1, Lnx-2 and Lnx-3 in their role as SDCs.<br><br>On Lnx-1, run the **`fdisk -l`** command and look for a ScaleIO device.<br><br>It should appear as **`/dev/scinia`**:<br><br>`Disk /dev/scinia: 17.2 GB, 17179869184 bytes`<br>`64 heads, 32 sectors/track, 16384 cylinders`<br>`Units = cylinders of 2048 * 512 = 1048576 bytes`<br>`Sector size (logical/physical): 512 bytes / 512 bytes`<br>`I/O size (minimum/optimal): 512 bytes / 512 bytes`<br>`Disk identifier: 0x00000000`<br><br>`Disk /dev/scinia doesn't contain a valid partition table` |

| Step | Action |
|---|---|
| 4. | Run the following command to prepare the device for mounting:<br><br>**`fdisk /dev/scinia`**<br><br>Choose (**n**)ew, (**p**)rimary partition, number **1**, accept the defaults for first and last cylinders, and (**w**)rite it to disk.<br><br>`[root@Lnx-1 ~]# fdisk /dev/scinia`<br>`Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel`<br>`Building a new DOS disklabel with disk identifier 0x1a42340d.`<br>`Changes will remain in memory only, until you decide to write them.`<br>`After that, of course, the previous content won't be recoverable.`<br><br>`Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)`<br><br>`WARNING: DOS-compatible mode is deprecated. It's strongly recommended to`<br>`switch off the mode (command 'c') and change display units to`<br>`sectors (command 'u').`<br><br>`Command (m for help): n`<br>`Command action`<br>`e extended`<br>`p primary partition (1-4)`<br>`p`<br>`Partition number (1-4): 1`<br>`First cylinder (1-16384, default 1): <Enter>`<br>`Using default value 1`<br>`Last cylinder, +cylinders or +size{K,M,G} (1-16384, default 16384): <Enter>`<br>`Using default value 16384`<br><br>`Command (m for help): w`<br>`The partition table has been altered!`<br><br>`Calling ioctl() to re-read partition table.`<br>`Syncing disks.` |

| Step | Action |
|---|---|
| 5. | Format the partition:<br><br>**mkfs.ext4 /dev/scinia**<br><br>[root@Lnx-1 ~]# mkfs.ext4 /dev/scinia<br>mke2fs 1.41.12 (17-May-2010)<br>Filesystem label=<br>OS type: Linux<br>Block size=4096 (log=2)<br>Fragment size=4096 (log=2)<br>Stride=0 blocks, Stripe width=0 blocks<br>1048576 inodes, 4194304 blocks<br>209715 blocks (5.00%) reserved for the super user<br>First data block=0<br>Maximum filesystem blocks=4294967296<br>128 block groups<br>32768 blocks per group, 32768 fragments per group<br>8192 inodes per group<br>Superblock backups stored on blocks:<br>32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,<br>4096000<br><br>Writing inode tables: done<br>Creating journal (32768 blocks): done<br>Writing superblocks and filesystem accounting information: done<br><br>This filesystem will be automatically checked every 24 mounts or<br>180 days, whichever comes first. Use tune2fs -c or -i to override.<br>[root@Lnx-1 ~]# |
| 6. | Notice that the amount of “Provisioned” storage just increased slightly, from 64 to 64.8GB.<br><br> 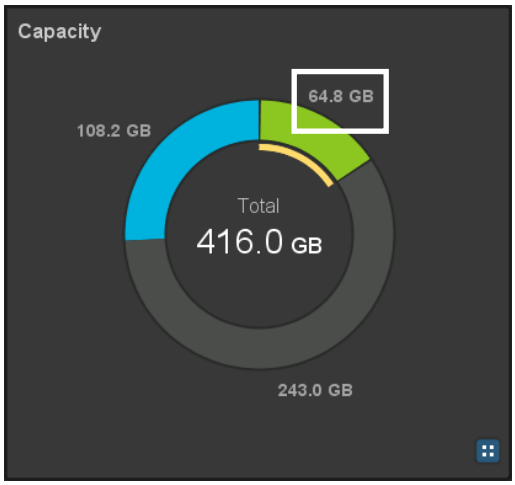 <br><br><br>The act of formatting involves writing some directories or tables to the disk device. Lnx-1 has a thin provisioned disk, so as data is written to previously-unused chunks of the device, ScaleIO will grow the volume to accommodate the writes. |

| Step | Action |
|---|---|
| 7. | Still on Lnx-1, create a mount point and mount the new device:<br><br>**mkdir -p /mnt/scaleio/1**<br><br>**mount /dev/scinia /mnt/scaleio/1**<br><br>Look for the new device:<br><br>**df -h**<br><br>[root@Lnx-1 ~]# mkdir -p /mnt/scaleio/1<br>[root@Lnx-1 ~]# mount /dev/scinia /mnt/scaleio/1<br>[root@Lnx-1 ~]# df -h<br>Filesystem Size Used Avail Use% Mounted on<br>/dev/mapper/vg_siorh62-lv_root<br>12G 3.4G 7.5G 32% /<br>tmpfs 1.9G 287M 1.6G 15% /dev/shm<br>/dev/sda1 485M 33M 427M 8% /boot<br>/dev/scinia 16G 172M 15G 2% /mnt/scaleio/1<br>[root@Lnx-1 ~]# |
| 8. | Connect to Lnx-2 and Lnx-3 and repeat Steps 4, 5 and 7 to map, format and mount vol 2 to Lnx-2 and vol 3 to Lnx-3 respectively |

| Step | Action |
|---|---|
| 9. | Recheck the dashboard in the GUI to see the change in capacity.<br> 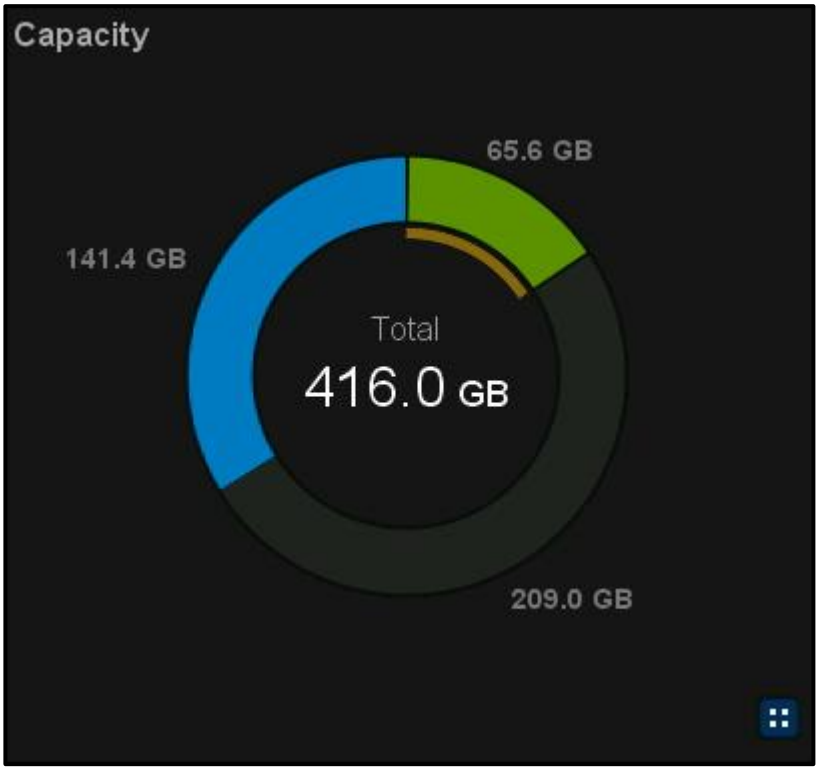 <br> |

## Lab 3-A: Task 4 – Map Volumes to Windows SDC

| Step | Action |
|---|---|
| | You will add a new Windows SDC, and map a volume for it. |
| 1. | We now have three Linux SDCs in the ScaleIO environment. Let's add a fourth SDC - but this time, a Windows server. |
| 2. | On the dc desktop, click on the win-2 icon to open a remote desktop connection (will automatically log on as administrator \| P@ssw0rd)<br> <br> |

| Step | Action |
|---|---|
| 3. | Open an explorer window to \\dc\share\sw\Windows. Double click the sdc installer. Complete the installation using default values.<br> 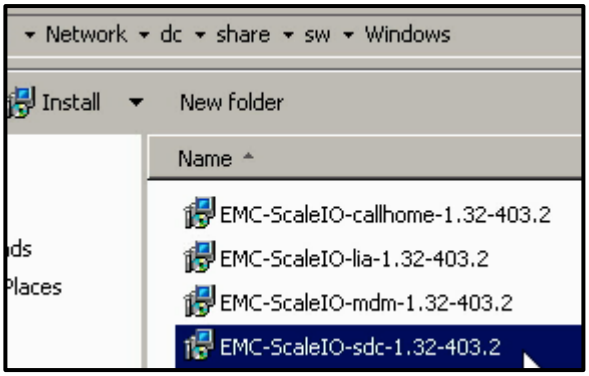 <br><br> <br> |
| 4. | Open a command prompt (Start → Command Promt)<br> <br> |
| 5. | Run the following commands to configure Win-2 as a client of our specific ScaleIO cluster:<br>**cd "C:\Program Files\EMC\ScaleIO\sdc\bin\"**<br>**drv_cfg.exe --add_mdm --ip 192.168.1.10**<br>**C:\Users\administrator.EMC>cd "C:\Program Files\EMC\ScaleIO\sdc\bin\"**<br>**C:\Program Files\EMC\ScaleIO\sdc\bin>drv_cfg.exe --add_mdm --ip 192.168.1.10**<br>**Calling kernel module to connect to MDM (192.168.1.10)**<br>**C:\Program Files\EMC\ScaleIO\sdc\bin>_** |

| Step | Action |
|---|---|
| 6. | Back on Lnx-1 (minimize the remote desktop connection to win-2, open putty window to lnx-1), a query to the MDM should now show the new (fourth) SDC.<br><br>`[root@Lnx-1 ~]# scli --query_all_sdc`<br>`MDM restricted SDC mode: Disabled`<br>`Query all SDC returned 4 SDC nodes.`<br>`SDC ID: c842438d00000000 Name: SDC_Lnx1 IP: 192.168.1.10`<br>`214-0EACBA3939F1`<br>`Read bandwidth: 0 IOPS 0 Bytes per-second`<br>`Write bandwidth: 0 IOPS 0 Bytes per-second`<br>`SDC ID: c842438e00000001 Name: SDC_Lnx3 IP: 192.168.1.12`<br>`B0E-5D6A4CB38FAF`<br>`Read bandwidth: 0 IOPS 0 Bytes per-second`<br>`Write bandwidth: 0 IOPS 0 Bytes per-second`<br>`SDC ID: c842438f00000002 Name: SDC_Lnx2 IP: 192.168.1.11`<br>`94C-6EDE4BA577F3`<br>`Read bandwidth: 0 IOPS 0 Bytes per-second`<br>`Write bandwidth: 0 IOPS 0 Bytes per-second`<br>`SDC ID: c8426a9d00000003 Name: N/A IP: 192.168.1.21 State`<br>`786AACE7E52`<br>`Read bandwidth: 0 IOPS 0 Bytes per-second`<br>`Write bandwidth: 0 IOPS 0 Bytes per-second` |
| 7. | While still connected to lnx-1, rename the new SDC.<br><br>`scli --rename_sdc --sdc_id` <SCD ID from above> `--new_name SDC_Win2`<br><br>`SDC ID: c8426a9d00000003 Name: N/A IP: 192.168.1.21 State: Connected GUID: EB630`<br>`786AACE7E52`<br>`Read bandwidth: 0 IOPS 0 Bytes per-second`<br>`Write bandwidth: 0 IOPS 0 Bytes per-second`<br>`[root@Lnx-1 ~]# scli --rename_sdc --sdc_id c8426a9d00000003 --new_name SDC_Win2`<br>`Successfully renamed sdc with ID c8426a9d00000003 to SDC_Win2`<br>`[root@Lnx-1 ~]#` |
| 8. | Map vol4 to SDC_Win2<br><br>`scli --map_volume_to_sdc --volume_name vol4 --sdc_name SDC_Win2`<br><br>`[root@Lnx-1 ~]# scli --map_volume_to_sdc --volume_name vol4 --sdc_name SDC_Win2`<br>`Successfully mapped volume vol4 to SDC SDC_Win2`<br>`[root@Lnx-1 ~]#` |

| Step | Action |
|---|---|
| 9. | On Win-2, open **Server Manager** and select **Storage → Disk Management**. If you get an "Initialize Disk" pop up, just cancel it.<br><br>The ScaleIO disk shows up ready to be initialized and formatted.<br><br> |

## Lab 3-A: Task 5 – Snapshot Volumes

| Step | Action |
|---|---|
| | You will create a snapshot for an existing volume, and map it from another sdc host. |
| 1. | First you will add "data" to a volume.<br><br>From lnx-1 shell, change to the mounted volume<br><br>`cd /mnt/scaleio/1`<br><br>`[root@Lnx-1 ~]# cd /mnt/scaleio/1` |
| 2. | Add a file to that directory. And confirm its presence.<br><br>`touch original`<br><br>`ls /mnt/scaleio/1`<br><br>`[root@Lnx-1 1]# touch original`<br>`[root@Lnx-1 1]# ls /mnt/scaleio/1`<br>`lost+found original`<br>`[root@Lnx-1 1]#` |
| 3. | Type the following command to ensure that any cached writes are written to the volume<br><br>`sync;sync`<br><br>`[root@Lnx-1 1]# sync;sync`<br>`[root@Lnx-1 1]#` |

| Step | Action |
|---|---|
| 4. | Create a snapshot of the volume called “vol1_snap”<br><br>**scli --snapshot_volume --volume_name vol1 --snapshot_name vol1_snap**<br><br>[root@Lnx-1 1]# scli --snapshot_volume --volume_name vol1 --snapshot_name vol1_snap<br>Snapshot created successfully<br>Source volume vol1 => b8fb434d00000004 vol1_snap |
| 5. | Map the volume to lnx-2<br><br>**scli --map_volume_to_sdc --volume_name vol1_snap --sdc_ip 192.168.1.11**<br><br>[root@Lnx-1 1]# scli --map_volume_to_sdc --volume_name vol1_snap --sdc_ip 192.168.1.11<br>Successfully mapped volume vol1_snap to SDC 192.168.1.11 |
| 6. | Connect to lnx-2 and locate the volume<br><br>**fdisk -l**<br><br>Disk /dev/scinib: 17.2 GB, 17179869184 bytes<br>64 heads, 32 sectors/track, 16384 cylinders<br>Units = cylinders of 2048 * 512 = 1048576 bytes<br>Sector size (logical/physical): 512 bytes / 512 bytes<br>I/O size (minimum/optimal): 512 bytes / 512 bytes<br>Disk identifier: 0x00000000<br><br>Disk /dev/scinib doesn't contain a valid partition table |
| 7. | Mount the volume.<br><br>**mkdir -p /mnt/scaleio/2**<br><br>**mount /dev/scinib /mnt/scaleio/2**<br><br>[root@Lnx-2 1]# mkdir -p /mnt/scaleio/2<br>[root@Lnx-2 1]# mount /dev/scinib /mnt/scaleio/2 |

| Step | Action |
|---|---|
| 8. | Take a look at the contents and add a new file<br><br>**cd /mnt/scaleio/2**<br><br>**ls**<br><br>**touch new-lnx-2**<br><br>**ls**<br><br>[root@Lnx-2 1]# cd /mnt/scaleio/2<br>[root@Lnx-2 2]# ls<br>lost+found original<br>[root@Lnx-2 2]# touch new-lnx-2<br>[root@Lnx-2 2]# ls<br>lost+found new-lnx-2 original |
| 9. | Unmap the volume<br><br>**cd /**<br><br>**Unmount /mnt/scaleio/2** |
| 10. | Return to lnx-1, unmap the snapshot volume from lnx-2, and remove the snapshot<br><br>**scli --unmap_volume_from_sdc --volume_name vol1_snap --sdc_name SDC_Lnx2**<br><br>**scli --remove_volume --volume_name vol1_snap**<br><br>[root@Lnx-1 1]# scli --unmap_volume_from_sdc --volume_name vol1_snap --sdc_name SDC_Lnx2<br>Volume will not be accessible to the SDC. Press 'y' to confirm.y<br>Successfully un-mapped volume vol1_snap from SDC SDC_Lnx2<br>[root@Lnx-1 1]# scli --remove_volume --volume_name vol1_snap<br>Removing a volume will erase all data on this volume. Press 'y' to confirm.y<br>Successfully removed volume 'vol1_snap' |

## Lab 3-A: Task 6 – Consistency Groups

| Step | Action |
|---|---|
| | You will create a consistency group (a snapshot of multiple volumes created at the same instant), and mount them on another host. |

| Step | Action |
|---|---|
| 1. | While still connected to lnx-1, change to the source volume directory, Add a new file and examine the contents of the volume<br><br>**cd /mnt/scaleio/1**<br><br>**touch consistent-data**<br><br>**ls**<br><br>[root@Lnx-1 1]# cd /mnt/scaleio/1<br>[root@Lnx-1 1]# touch consistent-data<br>[root@Lnx-1 1]# ls<br>consistent-data  lost+found  original |
| 2. | From lnx-2<br><br>Repeat the above step<br><br>[root@Lnx-2 1]# cd /mnt/scaleio/1<br>[root@Lnx-2 1]# touch consistent-data<br>[root@Lnx-2 1]# ls<br>consistent-data  lost+found  original |
| 3. | From lnx-1<br><br>Create a consistent snapshot of Lnx-1 vol1 and Lnx-2 vol2, called vol1_csnap and vol2_csnap respectively:<br><br>**scli --snapshot_volume --volume_name vol1,vol2 --snapshot_name vol1_csnap,vol2_csnap**<br><br>[root@Lnx-1 1]# scli --snapshot_volume --volume_name vol1,vol2 --snapshot_name vol1_csnap,vol2_csnap<br>Snapshots on 2 volumes created successfully<br>Consistency group ID: c359093700000002<br>Source volume vol1 => b8fb434e00000004 vol1_csnap<br>Source volume vol2 => b8fb434f00000005 vol2_csnap<br><br>Note the **Consistency Group ID** of your CG: _____ |

| Step | Action |
|---|---|
| 4. | From lnx-1<br><br>Map both consistent snapshots to lnx-3<br><br>**scli --map_volume_to_sdc --volume_name vol1_csnap -sdc_ip 192.168.1.12**<br><br>**scli --map_volume_to_sdc --volume_name vol2_csnap -sdc_ip 192.168.1.12**<br><br>[root@Lnx-1 1]# scli --map_volume_to_sdc --volume_name vol1_csnap --sdc_ip 192.168.1.12<br>Successfully mapped volume vol1_csnap to SDC 192.168.1.12<br>[root@Lnx-1 1]# scli --map_volume_to_sdc --volume_name vol2_csnap --sdc_ip 192.168.1.12<br>Successfully mapped volume vol2_csnap to SDC 192.168.1.12<br>[root@Lnx-1 1]# |
| 5. | From lnx-3<br><br>Examine the newly mapped volumes<br><br>**fdisk -l**<br><br>Disk /dev/scinib: 17.2 GB, 17179869184 bytes<br>64 heads, 32 sectors/track, 16384 cylinders<br>Units = cylinders of 2048 * 512 = 1048576 bytes<br>Sector size (logical/physical): 512 bytes / 512 bytes<br>I/O size (minimum/optimal): 512 bytes / 512 bytes<br>Disk identifier: 0x00000000<br><br>Disk /dev/scinib doesn't contain a valid partition table<br><br>Disk /dev/scinic: 17.2 GB, 17179869184 bytes<br>64 heads, 32 sectors/track, 16384 cylinders<br>Units = cylinders of 2048 * 512 = 1048576 bytes<br>Sector size (logical/physical): 512 bytes / 512 bytes<br>I/O size (minimum/optimal): 512 bytes / 512 bytes<br>Disk identifier: 0x00000000 |

| Step | Action |
|---|---|
| 6. | From lnx-3<br><br>Mount the volumes<br><br>**mkdir /mnt/scaleio/csnap1**<br><br>**mkdir /mnt/scaleio/csnap2**<br><br>**mount /dev/scinib /mnt/scaleio/csnap1**<br><br>**mount /dev/scinic /mnt/scaleio/csnap2**<br><br>`[root@Lnx-3 /]# mkdir /mnt/scaleio/csnap1`<br>`[root@Lnx-3 /]# mkdir /mnt/scaleio/csnap2`<br>`[root@Lnx-3 /]# mount /dev/scinib /mnt/scaleio/csnap1`<br>`[root@Lnx-3 /]# mount /dev/scinic /mnt/scaleio/csnap2` |
| 7. | From lnx-3<br><br>Confirm the contents of the volumes<br><br>**ls /mnt/scaleio/csnap1**<br><br>**ls /mnt/scaleio/csnap2**<br><br>`[root@Lnx-3 /]# ls /mnt/scaleio/csnap1`<br>`consistent-data  lost+found  original`<br>`[root@Lnx-3 /]# ls /mnt/scaleio/csnap2`<br>`consistent-data  lost+found  original` |
| 8. | From lnx-1<br><br>Unmap the consistent snapshots<br><br>**scli --unmap_volume_from_sdc --volume_name vol1_csnap --sdc_ip 192.168.1.12**<br><br>**scli --unmap_volume_from_sdc --volume_name vol2_csnap --sdc_ip 192.168.1.12**<br><br>`[root@Lnx-1 1]# scli --unmap_volume_from_sdc --volume_name vol1_csnap --sdc_ip 192.168.1.12`<br>`Volume will not be accessible to the SDC. Press 'y' to confirm.y`<br>`[root@Lnx-1 1]# scli --unmap_volume_from_sdc --volume_name vol1_csnap --sdc_ip 192.168.1.12`<br>`Volume will not be accessible to the SDC. Press 'y' to confirm.y` |

| Step | Action |
|---|---|
| 9. | From lnx-1<br><br>Remove the consistency group snapshots. First identify the consistency group ID<br><br>**scli --query_volume_tree --volume_name vol1**<br><br>[root@Lnx-1 1]# scli --query_volume_tree --volume_name vol1<br>VTree ID: 87e49fed00000000<br>Total snapshots size: 4.0 MB (4096 KB)<br>Total capacity in use: 399.0 MB (408576 KB)<br><br>>> Volume ID: b8fb1c3900000000 Name: vol1 Thin-provisioned Size: 16.0 GB<br>Capacity in use: 395.0 MB (404480 KB)<br>Trimmed capacity: 0 Bytes<br>Storage Pool a6041e4d00000000 Name: Cap_Pool<br>Protection Domain 18cafc9300000000 Name: Domain1<br>Creation time: 2015-07-30 10:41:39<br>Obfuscation is disabled<br>Uses RAM Read Cache<br>Read bandwidth: 0 IOPS 0 Bytes per-second<br>Write bandwidth: 0 IOPS 0 Bytes per-second<br>Mapped SDC:<br>SDC ID: c842438d00000000 IP: 192.168.1.10 Name: SDC_Lnx1<br>>> Volume ID: b8fb434e00000004 Name: vol1_csnap Snapshot of vol1 Size:<br>Storage Pool a6041e4d00000000 Name: Cap_Pool<br>Protection Domain 18cafc9300000000 Name: Domain1<br>Creation time: 2015-08-04 10:00:15<br>Obfuscation is disabled<br>Doesn't use RAM Read Cache<br>Read bandwidth: 0 IOPS 0 Bytes per-second<br>Write bandwidth: 0 IOPS 0 Bytes per-second<br>Consistency Group ID: c359093700000002 Mapped SDC:<br>SDC ID: c842438e00000001 IP: 192.168.1.12 Name: SDC_Lnx3 |
| 10. | From lnx-1<br><br>Next, remove the consistency group<br><br>**scli --remove_consistency_group_snapshots --consistency_group_id <cg ID>**<br><br>[root@Lnx-1 1]# scli --remove_consistency_group_snapshots --consistency_group_id c359093700000002<br>The command will remove a group of snapshots. Press 'y' to confirm. |

## Lab 3-A: Task 7 – Inactivation of a Protection Domain

| Step | Action |
|---|---|
| | You will deactivate a protection domain. |
| 1. | On Lnx-1, change to the home directory of the **root** account:<br><br>**`cd /root`** |
| 2. | Ensure that none of the file system contents are in the host file system cache. Unmount and mount the file system once:<br><br>**`umount /mnt/scaleio/1`**<br><br>**`mount /dev/scinia  /mnt/scaleio/1`**<br><br>**`df –h`**<br><br>At this point, we can be sure that any subsequent I/O to files in this file system will need to be serviced from the ScaleIO device – the host file system cache is empty. |
| 3. | Verify the status of the Protection Domain:<br><br>**`scli --query_protection_domain --protection_domain_name default |grep state`**<br><br>**Note:** Use grep to limit the output.<br><br>`[root@Lnx-1 1]# scli --query_protection_domain --protection_domain_name default |grep state`<br>`Operational state is Active`<br>`  4 SDS nodes have membership state 'Joined'`<br>`  4 SDS nodes have connection state 'Connected'`<br>`[root@Lnx-1 1]#` |

| Step | Action |
|---|---|
| 4. | Inactivate the Protection Domain:<br><br>**`scli --inactivate_protection_domain --protection_domain_name default`**<br><br>`[root@Lnx-1 1]# scli --inactivate_protection_domain --protection_domain_name default`<br>`All application IOs will not be serviced. Press 'y' to confirm that you want to proceed ?`<br>`Protection Domain default is being shutdown asynchronously`<br>`[root@Lnx-1 1]#`<br><br>Confirm that the current state of the Protection Domain is being reported as **inactive**, before proceeding:<br><br>**`scli --query_protection_domain --protection_domain_name default \|grep state`**<br><br>`[root@Lnx-1 /]# scli --query_protection_domain --protection_domain_name default \|grep state`<br>`Operational state is Inactive`<br>`  4 SDS nodes have membership state 'Joined'`<br>`  4 SDS nodes have connection state 'Connected'`<br>`[root@Lnx-1 /]#` |
| 5. | Check if we still have access to files on the mounted file system:<br><br>**`df -h`**<br><br>**`ls /mnt/scaleio/1`**<br><br>Did the **`ls`** command fail with an error? Is this expected behavior?<br><br>_____ |
| 6. | Activate the Protection Domain:<br><br>**`scli --activate_protection_domain --protection_domain_name default`**<br><br>Confirm that the current state of the Protection Domain is being reported as **active**, before proceeding:<br><br>**`scli --query_protection_domain --protection_domain_name default \| grep state`** |
| 7. | You should now have full read/write access to the file system. Verify this:<br><br>**`df -h`**<br><br>**`ls /mnt/scaleio/1`** |

## Lab 3-A: Task 8 – Rename the Protection Domain

| Step | Action |
|---|---|
| | You will rename the protection domain within the GUI |
| 1. | Back on Lnx-1, run the following command:<br><br>**`scli --query_all |grep Protection`**<br><br>`[root@Lnx-1 logs]# scli --query_all |grep Protection`<br>`Query all returned 1 Protection Domain:`<br>`Protection Domain default (Id: d17080e900000000) has 1 storage pools,`<br>`[root@Lnx-1 logs]#`<br><br>What is the Protection Domain name? _____ |
| 2. | On the dc desktop, within the ScaleIO GUI, go to the **Backend** view.<br><br> 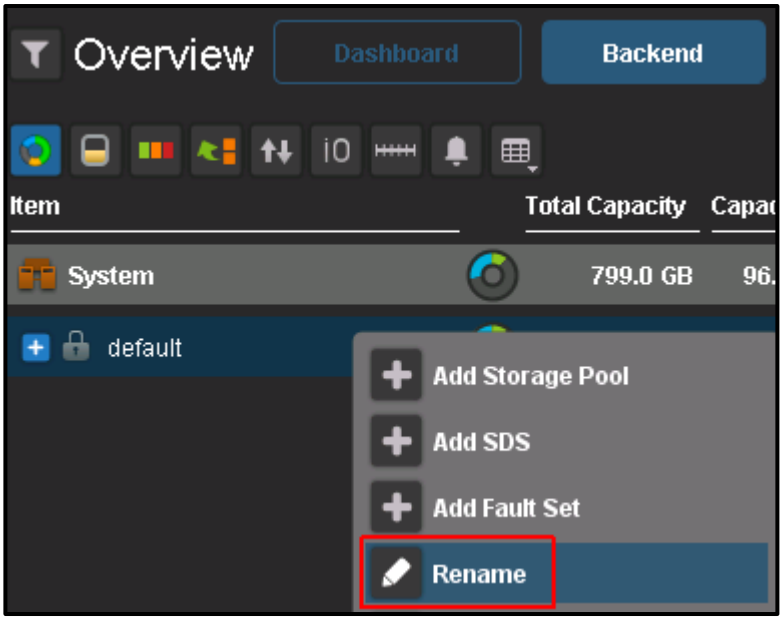 <br><br><br>Right-click on the Protection Domain, and select **Rename**. |

| Step | Action |
|---|---|
| 3. | In the **Rename Protection Domain** window, change the name of the Protection Domain to **Domain1** and click **OK**. Close when complete.<br> 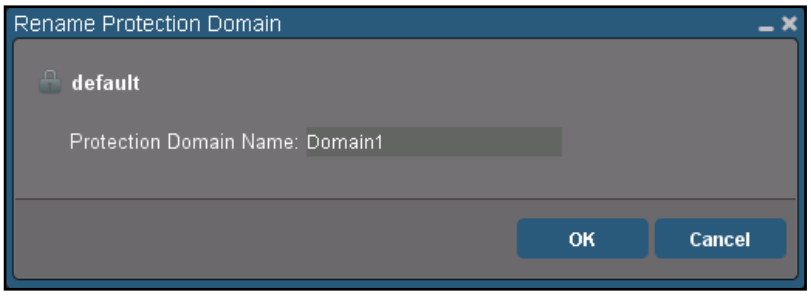 <br><br>Back on Lnx-1, run the following command:<br>**`scli --query_all \|grep Protection`**<br>`[root@Lnx-1 ~]# scli --query_all \|grep Protection`<br>`Query all returned 1 Protection Domain:`<br>`Protection Domain Domain1 (Id: b4ac019d00000000) has 1 storage pools,`<br>`72.0 GB (73728 MB) available for volume allocation`<br>`[root@Lnx-1 ~]#`<br>The name of the Protection Domain has changed to Domain1. |

## Lab 3-A: Task 9 - Rename Storage Pool

| Step | Action |
|---|---|
| | You will rename the Storage Pool within the GUI |
| 1. | Back on Lnx-1, run the following command:<br>**`scli --query_all \|grep Pool`**<br>`[root@Lnx-1 ~]# scli --query_all \|grep Pool`<br>`Storage Pool default (Id: 5ba3a08f00000000) has 4 volumes and 72.0 GB (73728 MB)`<br>`[root@Lnx-1 ~]#` |

| Step | Action |
|---|---|
| 2. | On the dc desktop, within the ScaleIO GUI, go to the **Backend** view.<br>Choose the drop-down in the upper-right-hand corner, selecting **By Storage Pools**.<br> 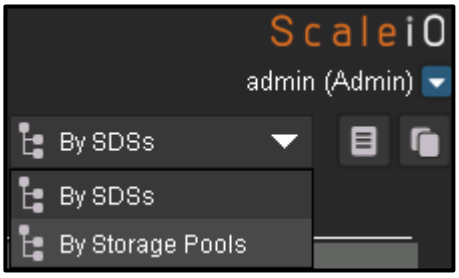 <br> |
| 3. | Underneath the Domain1 Protection Domain you will find the Storage Pool named “default”.<br> 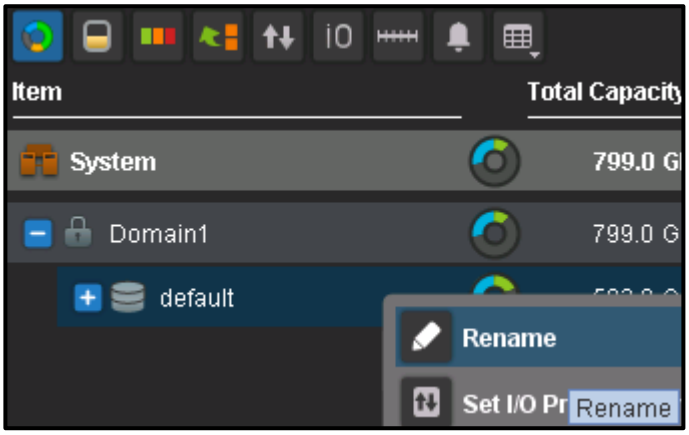 <br><br>Right-click on the Storage Pool and select **Rename**. |

| Step | Action |
|---|---|
| 4. | In the **Rename Storage Pool** window, change the name of the Storage Pool to **Cap_Pool** and click **OK**. Close when complete.<br><br> 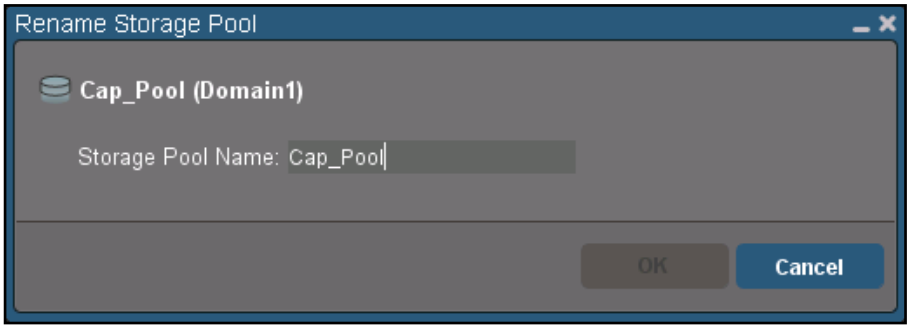 <br><br><br>Back on Lnx-1, run the following command:<br><br>**`scli --query_all \|grep Pool`**<br><br>`[root@Lnx-1 ~]# scli --query_all \|grep Pool`<br>`Storage Pool Cap_Pool (Id: 5ba3a08f00000000)`<br>`[root@Lnx-1 ~]#`<br><br>The name of the Storage Pool has changed to **Cap_Pool**. |

## Lab 3-A: Task 9 - Rename Storage Pool

| Step | Action |
|---|---|
| | You will rename the SDS systems within the GUI |
| 1. | Back on Lnx-1, run the following command:<br><br>**`scli --query_all_sds`**<br><br>`Query-all-SDS returned 4 SDS nodes.`<br><br>`Protection Domain d17080e900000000 Name: Domain1`<br>`SDS ID: e7f5951e00000003 Name: SDSWin2 IP: 192.168.1.21 Port: 7072`<br>`SDS ID: e7f5951d00000002 Name: SDS_[192.168.1.10] IP: 192.168.1.10 Port: 7072`<br>`SDS ID: e7f5951c00000001 Name: SDS_[192.168.1.11] IP: 192.168.1.11 Port: 7072`<br>`SDS ID: e7f5951b00000000 Name: SDS_[192.168.1.12] IP: 192.168.1.12 Port: 7072` |

| Step | Action |
|---|---|
| 2. | On the Win-1 desktop, within the ScaleIO GUI, go to the **Backend** view.<br><br> 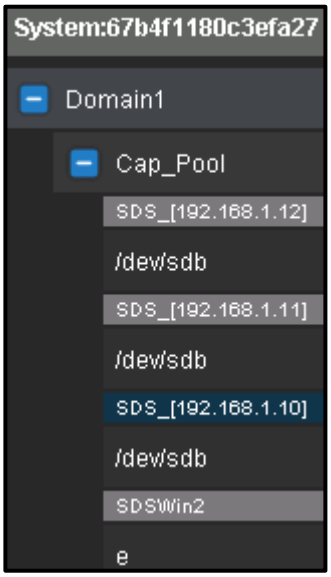 <br><br><br>Click on the row labeled **SDS_[192.168.1.10]**.<br>Right-click and select **rename**. |
| 3. | In the **Rename SDS** window, change the name of the SDS to **SDSLnx1** and click **OK**. Close when complete.<br><br> 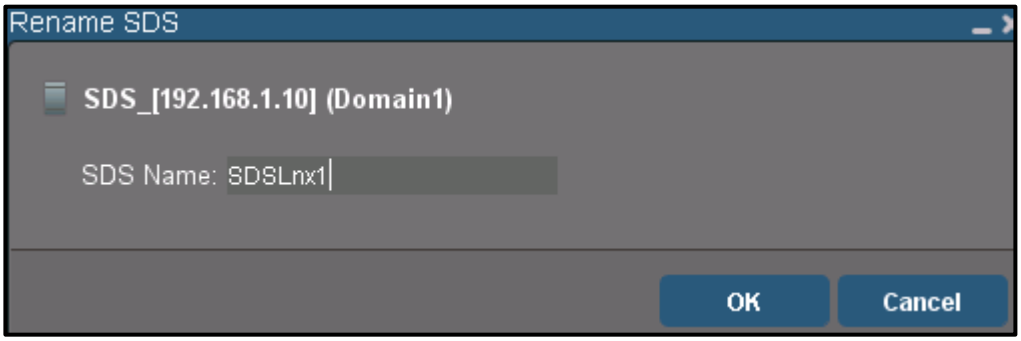 <br> |
| 4. | Repeat the steps above, renaming SDS 192.168.1.11 to **SDSLnx2** and SDS 192.168.1.12 to **SDSLnx3**. |

*End of Lab Exercise 3-A*

## Lab Exercise 3-B: ScaleIO Configuration and Management - ESXi Cluster

| | |
|---|---|
| **Purpose:** | Perform basic configuration, management and monitoring of a ScaleIO cluster in a VMware environment. |

| | |
|---|---|
| **Task:** | ✓ Perform operations that are supported by the ScaleIO vSphere plug-in beyond initial cluster deployment |

| | |
|---|---|
| **References:** | ScaleIO Lab Diagram (graphically showing relevant details for all VMs within a student pod – refer to Page 2) |

## Lab 3-B: Task 1 - Explore vSphere Plugin

| Step | Action |
|---|---|
| | Explore the basic information available through the ScaleIO vSphere plugin. |
| 1. | From the dc desktop, logon to the vSphere Web Client (desktop icon) as root \| P@ssw0rd 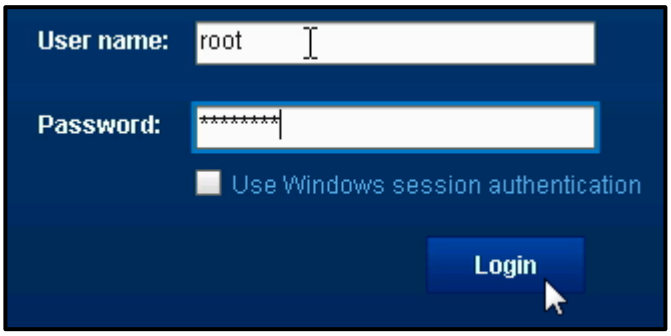  |
| 2. | On the Web Client Home Screen, click on the **EMC ScaleIO** icon. 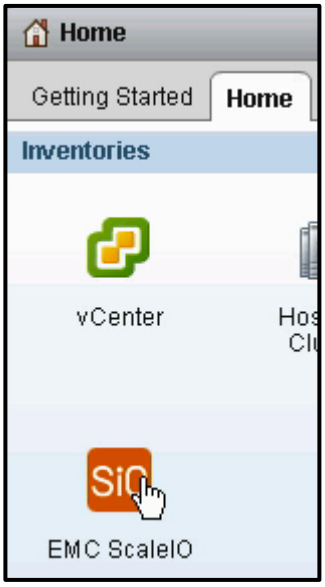  |

| Step | Action |
|---|---|
| 3. | On the lower-left corner, you can find the ScaleIO WebSphere plug-in options.<br> 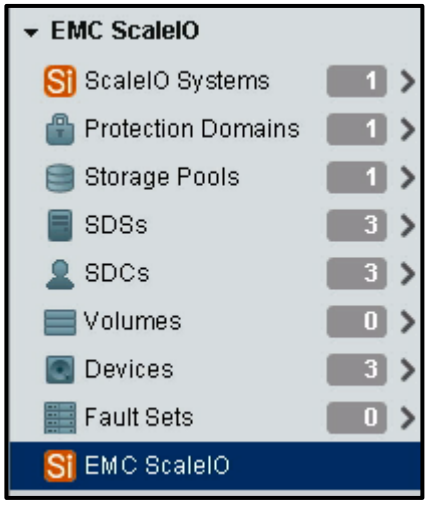 <br> |
| 4. | Take a moment to explore the views. After clicking and viewing each, you can return to the vCenter view by clicking the **vCenter** link in the upper-left corner of the screen.<br> 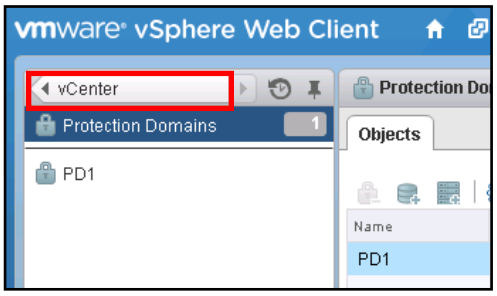 <br> |

## Lab 3-B: Task 2 - Create and Map ScaleIO Volumes

| Step | Action |
|---|---|
| | You will create volumes, and make them available to all the ESXi (SDC) hosts, within the ScaleIO plugin. |

| Step | Action |
|---|---|
| 1. | Navigate to the Storage Pool view and highlight the **SP1** Storage Pool. On the **Actions** drop-down menu, choose **Create volume**.<br> 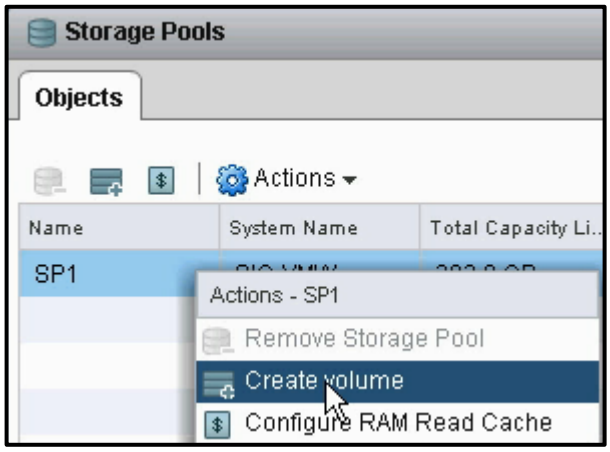 <br> |
| 2. | In the pop-up window that appears, enter the information to create a volume named **vol_1**, sized at **32** GB, **Thin** provisioned. Check **Map Volume to ESXs** and select **SIO**, which will automatically choose all three ESXi hosts within the ESXi cluster named **SIO**.<br> 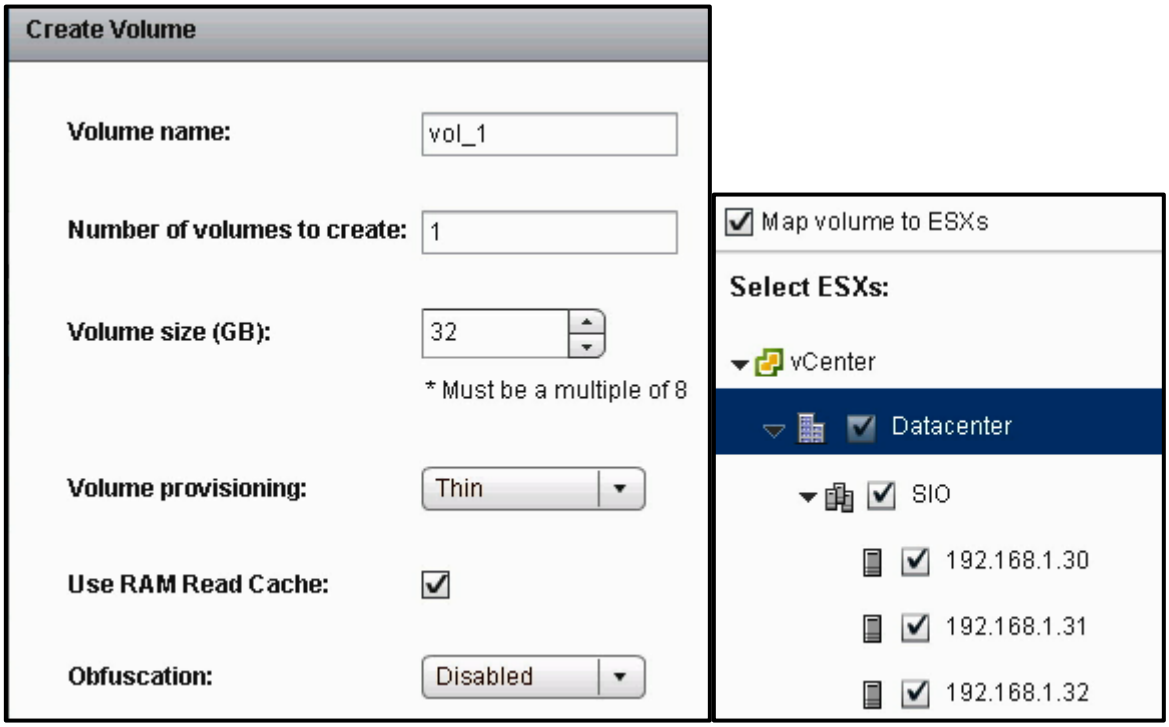 <br> |

| Step | Action |
|---|---|
| 3. | Click **OK**. If prompted, enter the ScaleIO admin password.<br> 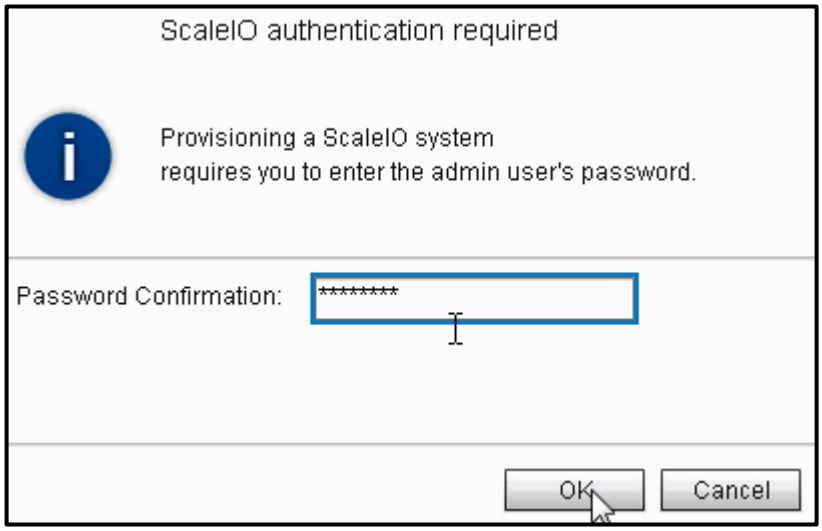 <br> |
| 4. | When finished, the following will appear:<br> 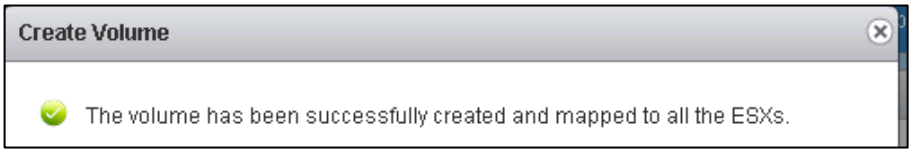 <br><br>Click **Close**. |
| 5. | Repeat the steps 1-4, creating **vol_2**, sized at **24** GB, **Thick** provisioned. Choose **Map Volume to ESXs** and select **SIO**, which will automatically choose all three ESXi hosts in the **SIO** cluster. |
| 6. | Go to the Volumes view.<br> 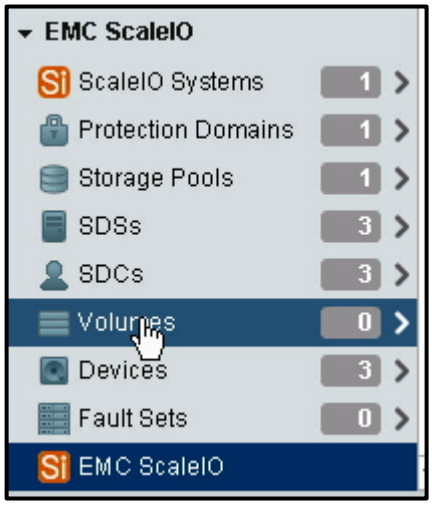 <br> |

| Step | Action |
|---|---|
| 7. | There, you should see your two newly-created volumes:<br> 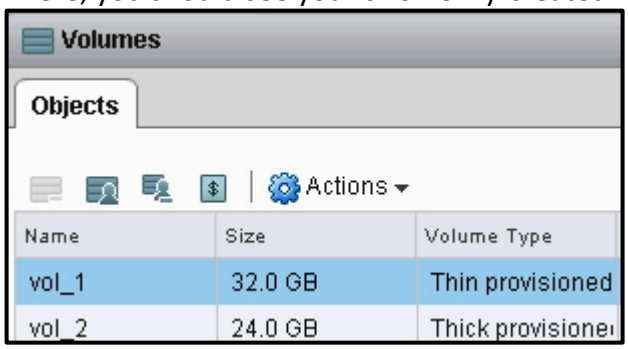 <br> |
| 8. | Navigate to vCenter Home → Hosts and Clusters<br> 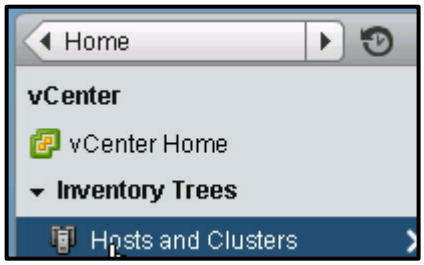 <br> |
| 9. | Select the 192.168. host<br> 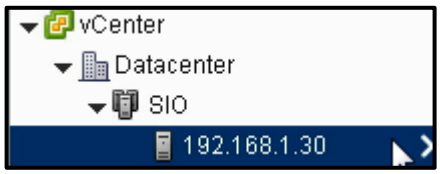 <br> |
| 10. | In the right pane, click Manage → Storage → Storage Devices<br> 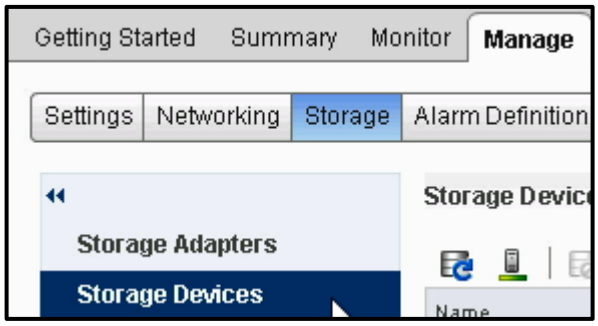 <br> |

| Step | Action |
|---|---|
| 11. | Notice the two new devices (easily identified by size)<br> 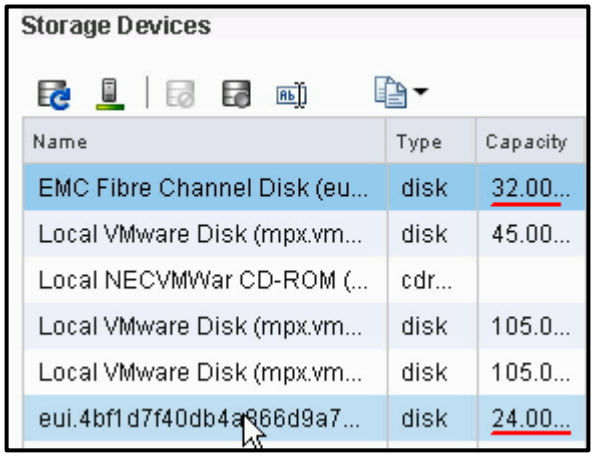 <br> |

## Lab 3-B: Task 3 - Create Datastores

| Step | Action |
|---|---|
| | You will create datastores on the volume just created and mapped. This is handled through the standard vSphere Web Client. |
| 1. | While in the same vSphere view, right-click on 192.168. and click New Datastore<br> 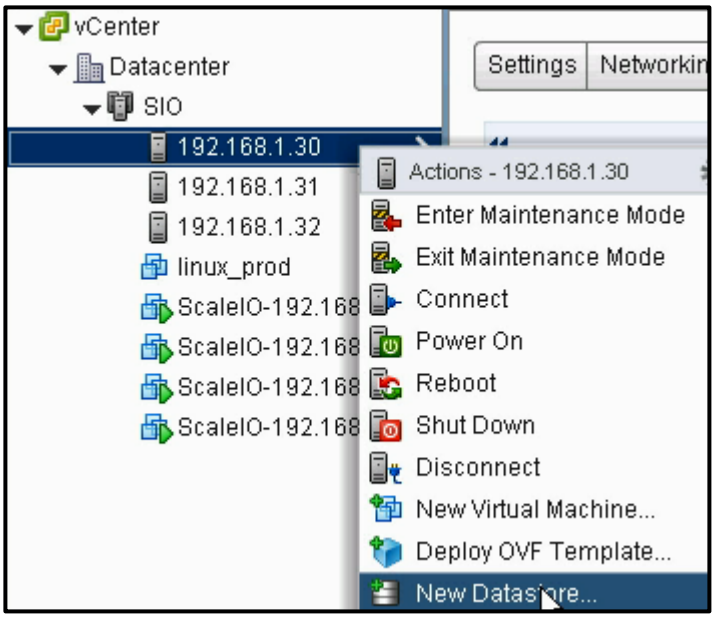 <br> |

| Step | Action |
|---|---|
| 2. | 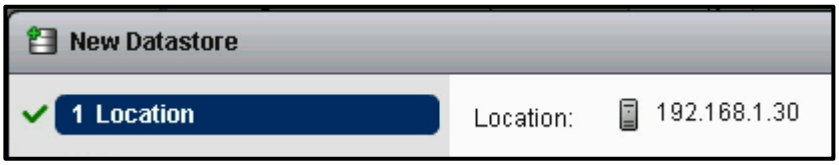 <br>Click **Next**.<br> 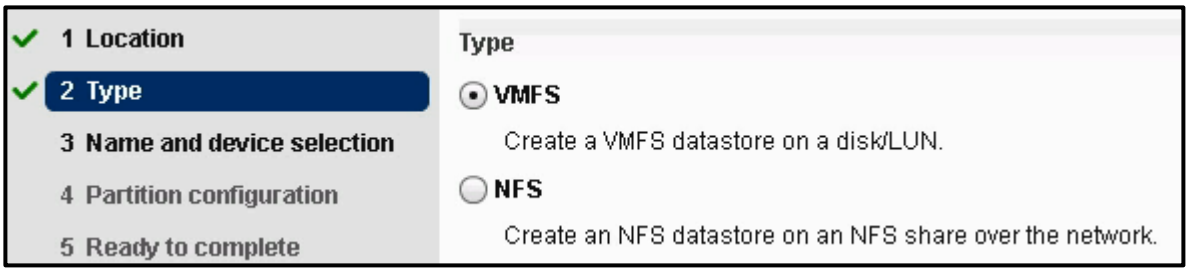 <br>Choose **VMFS** and click **Next**. |
| 3. | 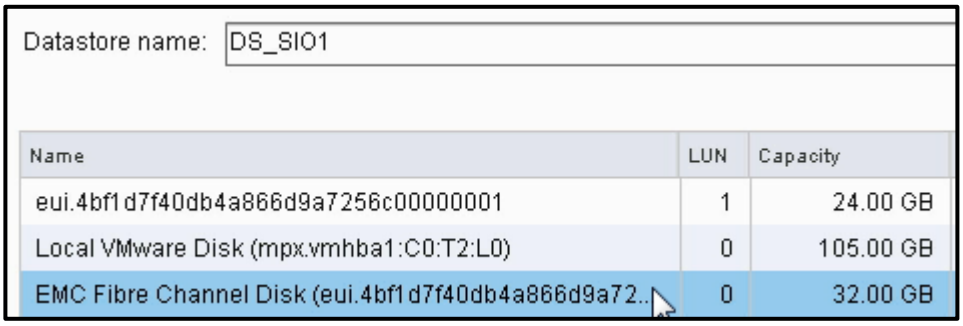 <br>Datastore name: **DS_SIO1**.<br>Choose the 32 GB iSCSI disk.<br>Click **Next**. |

| Step | Action |
|---|---|
| 4. | 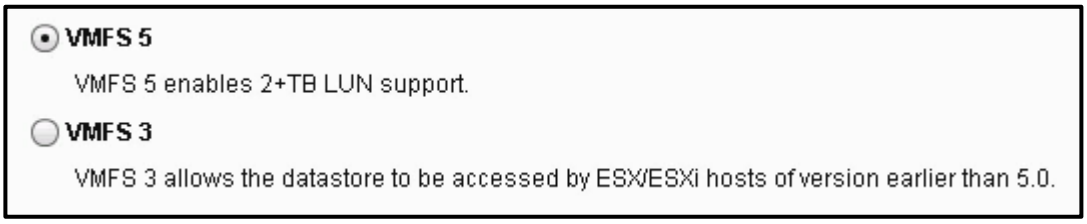 <br>Choose VMFS 5 and click **Next**. |
| 5. | On **Partition Configuration**, select **Use all available partitions**. Click **Next**.<br> 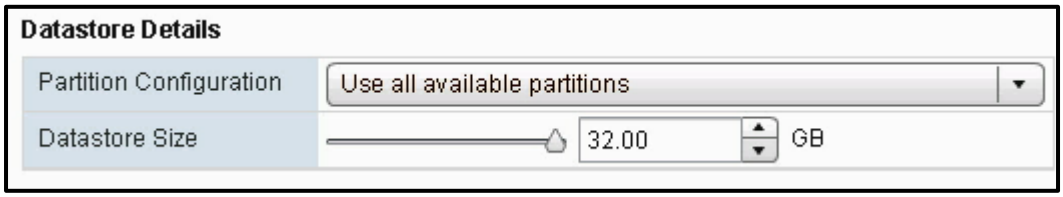 |
| 6. | 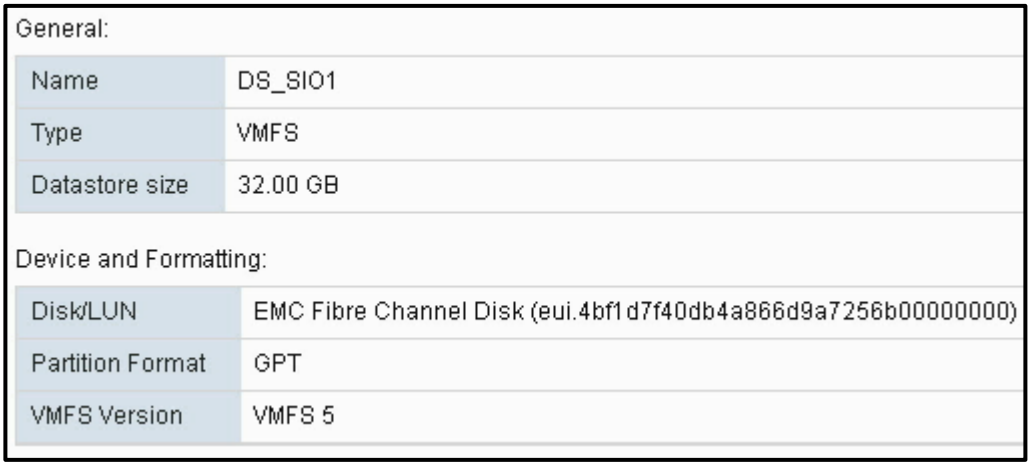 <br>Review and click **Finish**. |
| 7. | Repeat these steps to create Datastore DS_SIO2, using the 24 GB iSCSI disk. |

| Step | Action |
|---|---|
| 8. | At this point, all three ESXi hosts should be able to see Datastore DS_SIO1 and DS_SIO2 under the **Related Objects/Datastores** tab.<br> 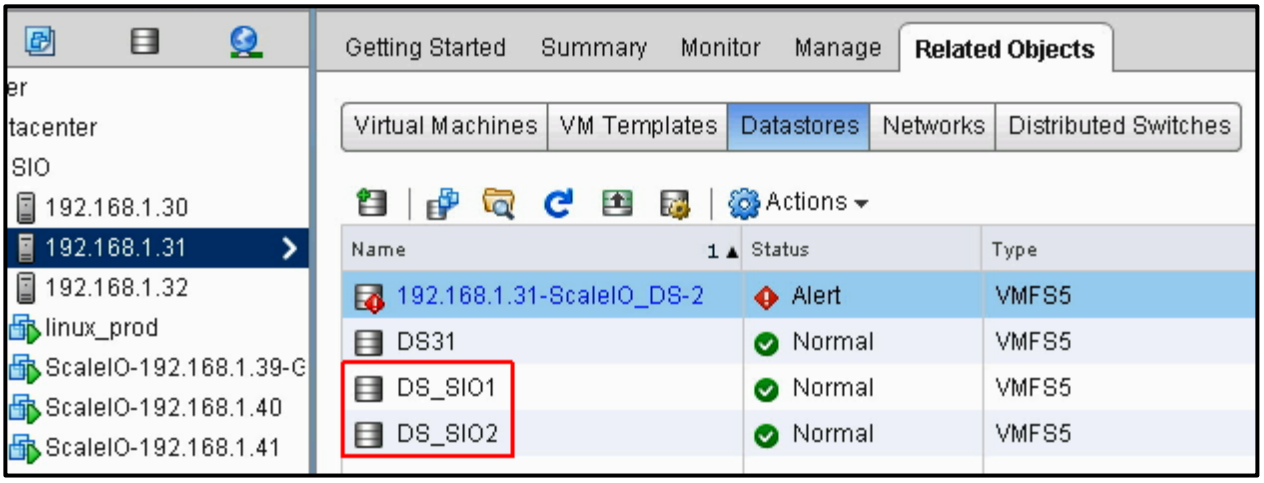 <br><br>*(the error on the top datastore relates to the VMDK creation in the initial deployment – it consumes nearly all the space on the datastore – this is expected)* |

## Lab 3-B: Task 4 - Migrate Live VM to ScaleIO Datastore

| Step | Action |
|---|---|
| | Now that you have a datastore available to all ESX hosts, you can migrate any VM to that datastore. This is done in the standard vSphere Web Client. |
| 1. | In vSphere → Hosts and Clusters, right-click on the linux-prod VM and click Power On<br> 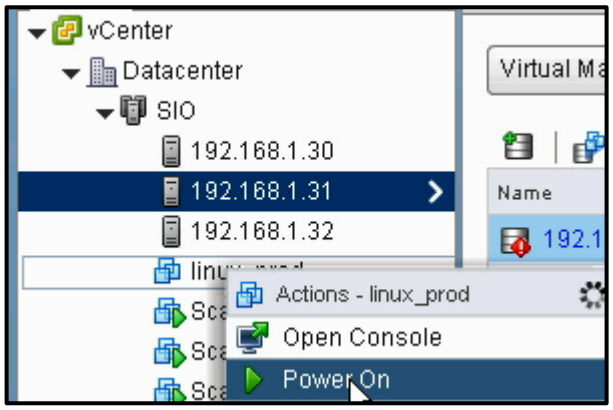 <br> |

| Step | Action |
|---|---|
| 2. | Connect to linux_prod via desktop icon<br> <br><br>(will log on automatically as root \| P@ssw0rd) |
| 3. | Examine the disk layout of the VM. See how many devices and how much space is being consumed.<br>**fdisk -l**<br>**df -h**<br>`linuxprod:~ # fdisk -l`<br>`Disk /dev/sda: 3221 MB, 3221225472 bytes`<br>`255 heads, 63 sectors/track, 391 cylinders, total 6291456 sectors`<br>`Units = sectors of 1 * 512 = 512 bytes`<br>`Sector size (logical/physical): 512 bytes / 512 bytes`<br>`I/O size (minimum/optimal): 512 bytes / 512 bytes`<br>`Disk identifier: 0x0006abda`<br>`Device Boot Start End Blocks Id System`<br>`/dev/sda1 2048 1028095 513024 82 Linux swap`<br>`/dev/sda2 * 1028096 6291455 2631680 83 Linux`<br>`linuxprod:~ # df -h`<br>`Filesystem Size Used Avail Use% Mounted on`<br>`/dev/sda2 2.5G 1.1G 1.3G 46% /`<br>`devtmpfs 940M 116K 940M 1% /dev`<br>`tmpfs 940M 0 940M 0% /dev/shm` |
| 4. | Discover the current VMware datastore hosting this VM. In vSphere, select the VM. Click Summary<br> 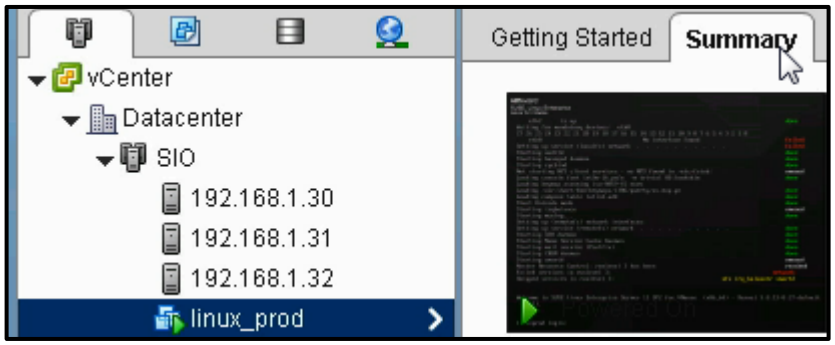 <br> |

| Step | Action |
|---|---|
| 5. | Under VM Hardware, expand the Hard Disk Section<br> 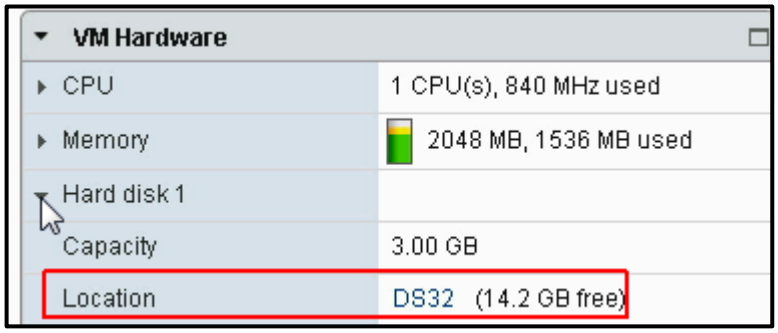 <br> |
| 6. | Back on vSphere, right-click the linux_prod VM and choose **Migrate**.<br> <br> |
| 7. | On the resulting pop-up windows, choose **Change Datastore** and click **Next**.<br> 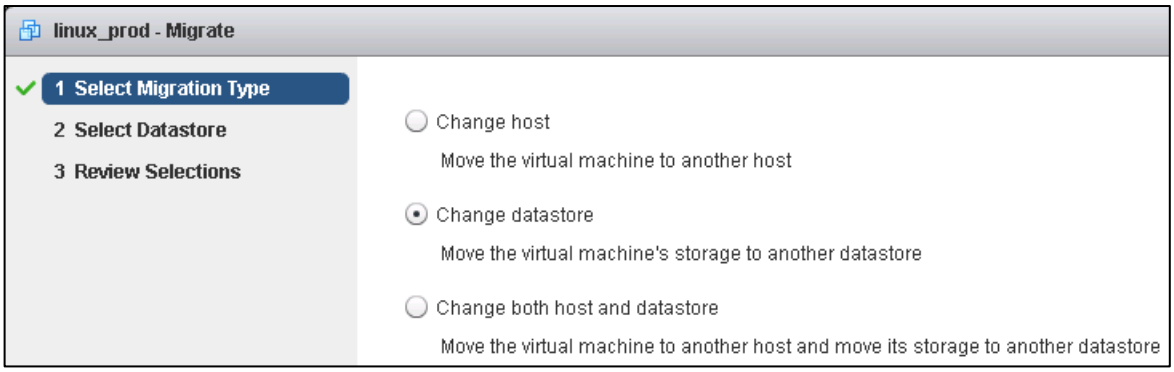 <br> |

| Step | Action |
|---|---|
| 8. | Select the **DS_SIO1** Datastore and click **Next**.<br> 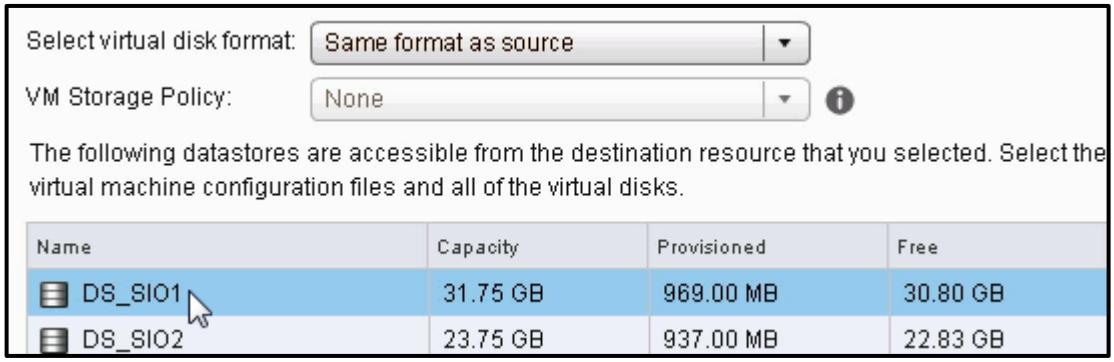 <br> |
| 9. | Review the choices displayed. If correct, click **Finish**.<br> 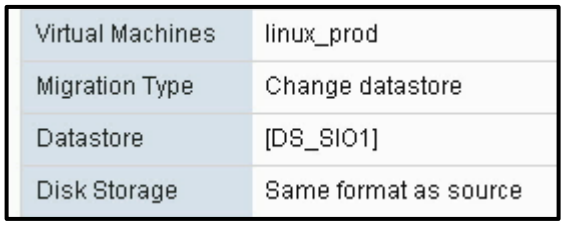 <br> |
| 10. | Wait for the migration to complete. We have just accomplished an online migration of a VM from a local datastore to a ScaleIO-based datastore that is accessible by all three ESX servers.<br> <br> |

## Lab 3-B: Task 5 - Migrate Live VM to other host

| Step | Action |
|---|---|
| | Now that the VM is hosted a datastore available to all ESX hosts, you can migrate any VM to any other ESX host that is attached to that datastore. This is done in the standard vSphere Web Client. |

| Step | Action |
|---|---|
| 1. | On the VMware Web Client, select the linux_prod VM. Its current host should be reported as 192.168.1.32.<br> 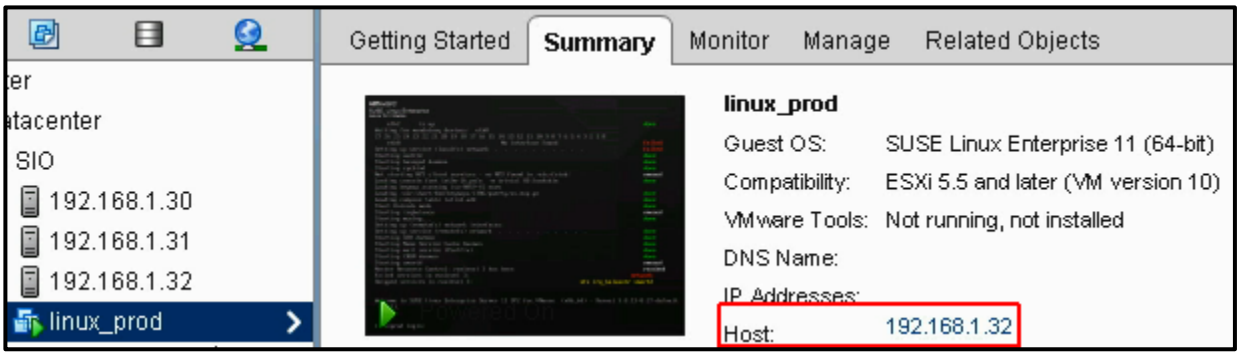 <br> |
| 2. | Right-click linux_prod and choose **Migrate…**.<br> <br> |
| 3. | On the first pop-up, leave the default (Change Host) and click **Next**.<br> 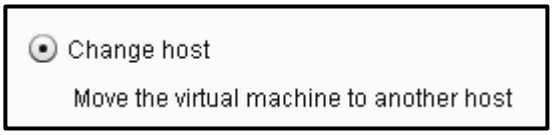 <br> |
| 4. | On the next screen, make sure that SIO is chosen and click **Next**.<br> 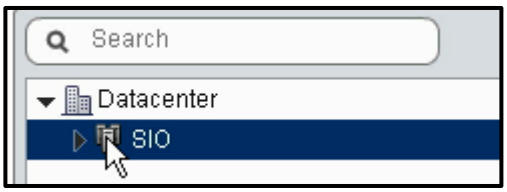 <br> |

| Step | Action |
|---|---|
| 5. | Choose the host 192.168.  and click **Next**.<br><br>Host Name / Status<br>192.168.1.30 — Normal<br>192.168.1.31 — Normal<br>192.168.1.32 — Normal |
| 6. | On the next screen, accept the defaults (Reserve CPU…), and click **Next**.<br><br>vMotion will move the virtual machine while it remains powered on.<br>Reserve CPU for optimal vMotion performance (Recommended)<br>If sufficient CPU resources are not immediately available, vMotion will not be initiated.<br>Perform with available CPU resources<br>If there is a lack of CPU resources, the duration of vMotion might be extended. |
| 7. | On the following screen, review the choices and, if correct, click **Finish**.<br><br>Virtual Machines: linux_prod<br>Migration Type: Change host<br>Resource Pool: Resources<br>Host: 192.168.1.30<br>VMotion Priority: Optimal vMotion with Reserved CPU |

| Step | Action |
|---|---|
| 8. | Once the migration completes, check that the **Summary** page for prod_linux correctly shows the new host it is running on: 192.168. . 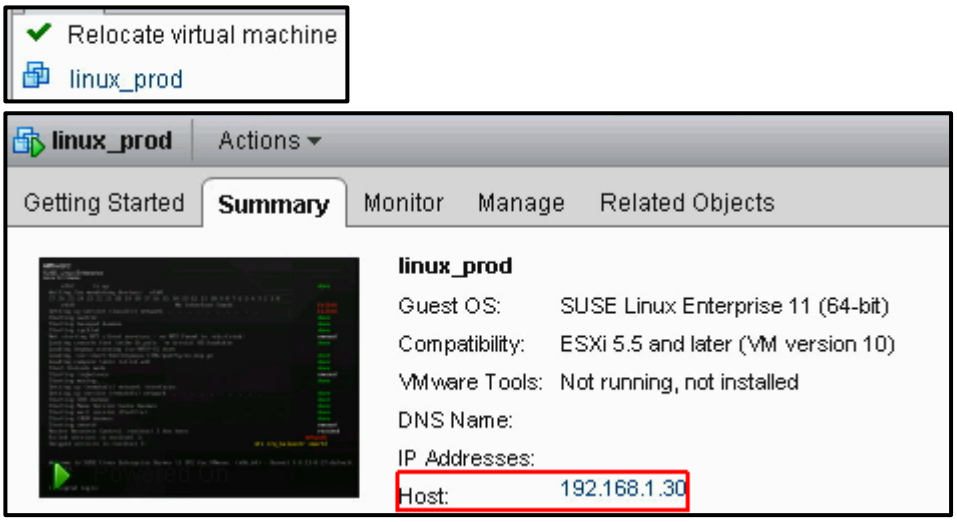  |

*End of Lab Exercise 3-B*

## Lab Exercise 4: Monitoring and Performance Tuning

| | |
|---|---|
| **Purpose:** | Perform various operations using the ScaleIO GUI and the CLI to monitor and tune a ScaleIO storage environment. |

| | |
|---|---|
| **Tasks:** | ✓ Observing Rebalance and Rebuild using Dashboard GUI and CLI<br>✓ Adjusting the Rebalance and Rebuild Policy<br>✓ Configure Quality of Service for a ScaleIO volume<br>✓ Explore and reconfigure the RAM cache setting<br>✓ View ScaleIO reported events and ScaleIO generated log files<br>✓ Redirect ScaleIO event log to a centralized remote syslog server |

| | |
|---|---|
| **References:** | ScaleIO Lab Diagram (graphically showing relevant details for all VMs within a student pod – refer to Page 2) |

## Lab 4: Task 1 - Monitoring, Rebalance and Rebuild When Adding/Removing an SDS

| Step | Action |
|---|---|
| | You will install and attach a new SDS to the cluster, then monitor the rebuild activity as data is redistributed to the new SDS. Then, you'll disconnect the new SDS and watch it rebuild the data from the surviving SDS nodes. |
| 1. | On dc, close and re-open open the ScaleIO GUI, logging into the Linux ScaleIO Cluster (192.168.1.10 or 192.168.1.11). Reminder: use **admin/P@ssw0rd**.<br> 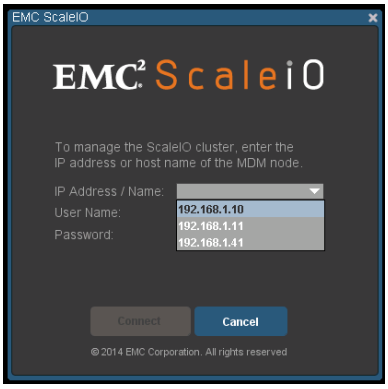 <br><br>Once logged in, notice that at this point we have 4 SDSs in the environment.<br> <br> |

| Step | Action |
|---|---|
| 2. | Go to the **Backend** view, sort by SDSs, and expand the Protection Domain.<br> 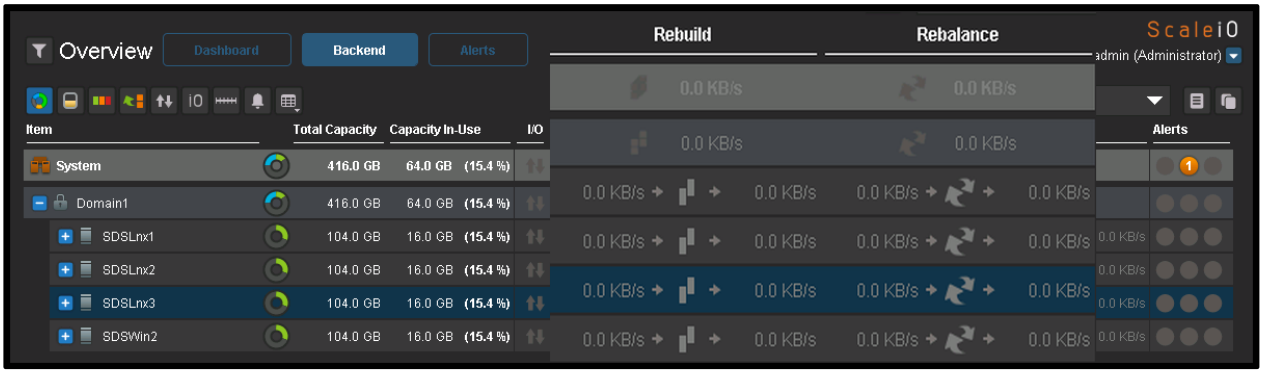 <br><br>Note that there is no Rebuild or Rebalance activity at the moment. |
| 3. | From dc, open a Remote Desktop Session to Win-3 (192.168.1.22).<br><br>**Hint:** Click **Start** and type the word “remote” in the box provided. Remote Desktop will appear among the available choices above the box. Log in as “administrator” and use the password “P@ssw0rd”. |
| 4. | On Win-3 go to **Server Manager→Storage→Disk Management**, confirm the presence of Disk 2 as a 105 GB disk device.<br> 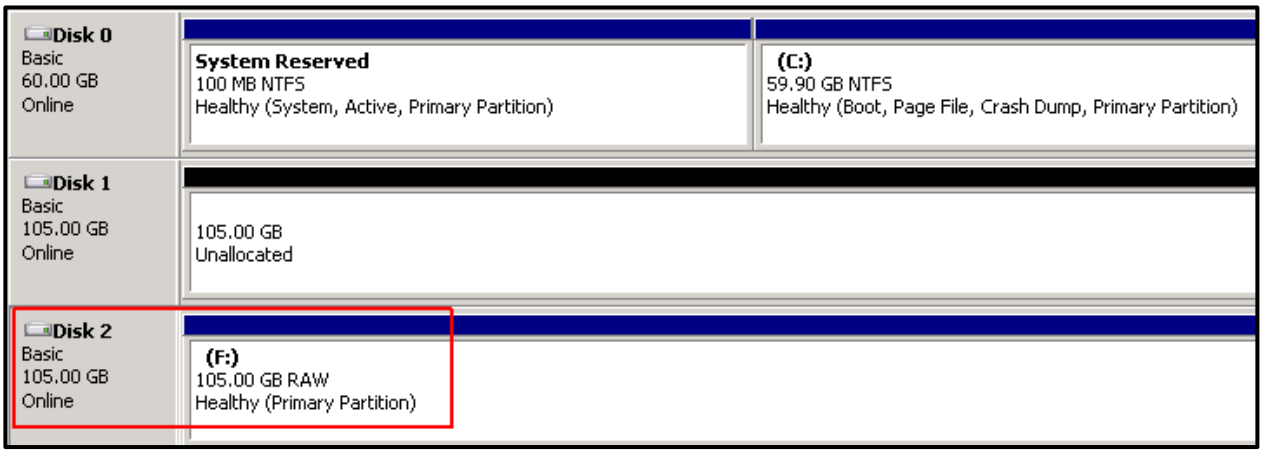 <br> |

| Step | Action |
|---|---|
| 5. | On the Win-3 server, right-click **Start** (lower-left-hand corner).<br><br>Select **Open Windows Explorer** to the **\\dc\share\sw\Windows** folder, and run the sds installer<br><br> 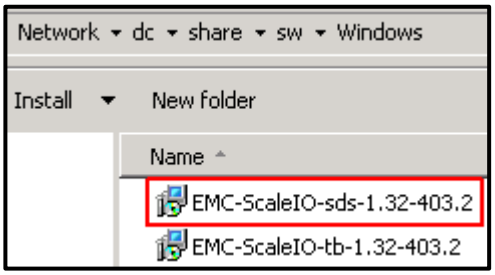 <br><br><br>Complete the installation using default values. |
| 6. | Back on the dc server, in the ScaleIO GUI go to the **Backend** view.<br><br>Right-click the Domain1 domain and select **Add SDS**.<br><br> 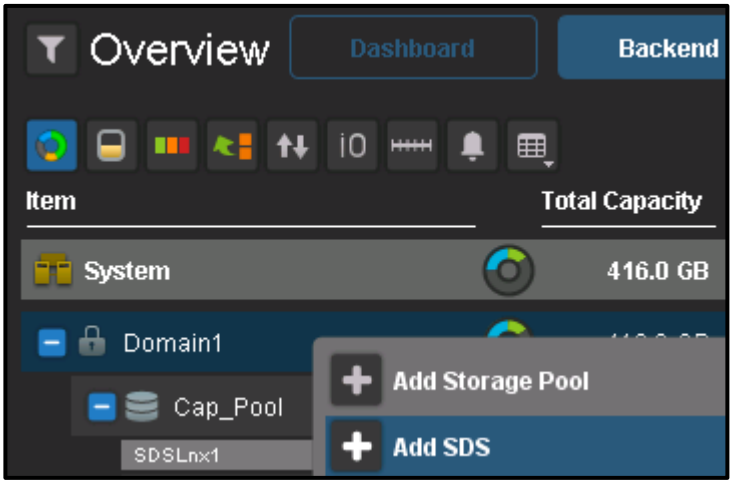 <br> |

| Step | Action |
|---|---|
| 7. | On the **Add SDS to Protection Domain** window, input the following:<br><br>Name: **SDSWin3**<br>IP Address: **192.168.1.22**<br>Path: **F**<br>Name: **Disk1**<br><br> 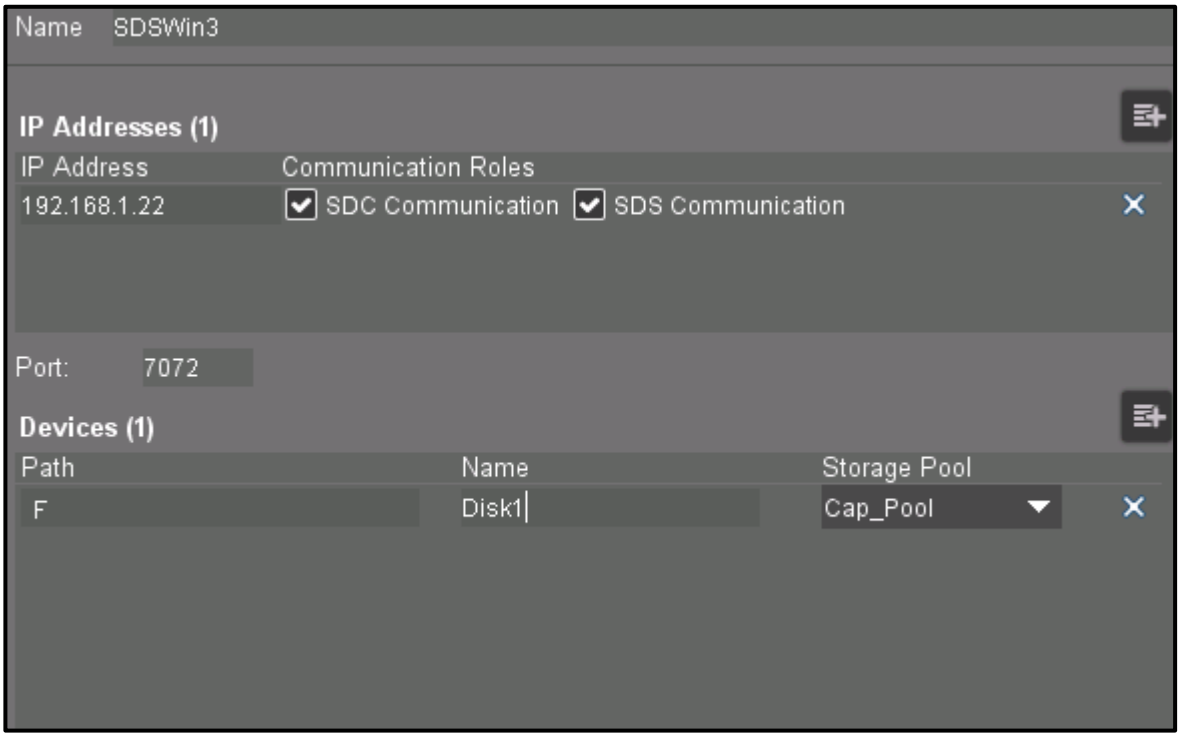 <br><br><br>Click **OK** when ready. |

| Step | Action |
|---|---|
| 8. | You will see Rebalance traffic, as the Domain adjusts to the presence of the new SDS. |

| Item | Total Capacity | Capacity In-Use | I/O Bandwidth | IOPS | Rebuild | Rebalance |
|---|---|---|---|---|---|---|
| System | 520.0 GB | 76.6 GB (14.7 %) | 0.0 KB/s | 0 | 0.0 KB/s | 61.2 MB/s |
| Domain1 | 520.0 GB | 76.6 GB (14.7 %) | 0.0 KB/s | 0 | 0.0 KB/s | 61.2 MB/s |
| SDS_[192.168.1.10] | 104.0 GB | 15.7 GB (15.1 %) | 0.0 KB/s | 0 | 0.0 KB/s → 0.0 KB/s | 0.0 KB/s → 18.4 MB/s |
| SDS_[192.168.1.11] | 104.0 GB | 15.6 GB (15.0 %) | 0.0 KB/s | 0 | 0.0 KB/s → 0.0 KB/s | 0.0 KB/s → 10.7 MB/s |
| SDS_[192.168.1.12] | 104.0 GB | 15.7 GB (15.1 %) | 0.0 KB/s | 0 | 0.0 KB/s → 0.0 KB/s | 0.0 KB/s → 9.3 MB/s |
| SDSWin2 | 104.0 GB | 16.0 GB (15.4 %) | 0.0 KB/s | 0 | 0.0 KB/s → 0.0 KB/s | 0.0 KB/s → 22.8 MB/s |
| SDSWin3 | 104.0 GB | 13.5 GB (13.0 %) | 0.0 KB/s | 0 | 0.0 KB/s → 0.0 KB/s | 60.4 MB/s → 0.0 KB/s |

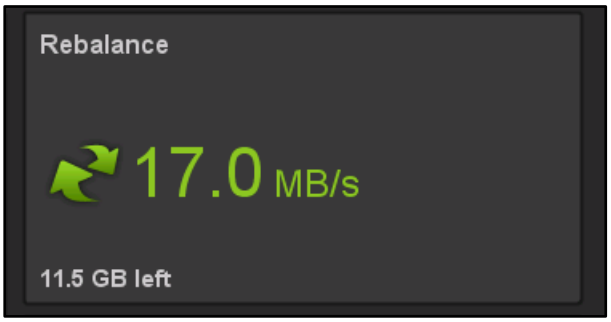

| Step | Action |
|---|---|
| 9. | Connect to lnx-1.<br><br>Run the following command:<br><br>**`scli --query_protection_domain --protection_domain_name Domain1 \| grep SDS`**<br><br>Take a moment to look over the SDS section, then run the command again, filtering output for the word “rebuild”, and then again, filtering for the word “rebalance”.<br><br>**`scli --query_protection_domain --protection_domain_name Domain1\|grep -i rebuild`**<br><br>**`scli --query_protection_domain --protection_domain_name Domain1\|grep -i rebalance`**<br><br>`[root@Lnx-1 ~]# scli --query_protection_domain --protection_domain_name Domain1\|grep -i rebuild`<br>`  The number of parallel rebuild/rebalance jobs: 2`<br>`  Rebuild is Enabled and using Limit-Concurrent-IO policy with the following parameters:`<br>`  The number of parallel rebuild/rebalance jobs: 2`<br>`  Rebuild is Enabled and using Limit-Concurrent-IO policy with the following parameters:`<br>`  0 Bytes active fwd-rebuild capacity`<br>`  0 Bytes pending fwd-rebuild capacity`<br>`  0 Bytes active bck-rebuild capacity`<br>`  0 Bytes pending bck-rebuild capacity`<br>`  0 Bytes fwd-rebuild capacity`<br>`  0 Bytes bck-rebuild capacity`<br>`  Backward-rebuild-reads  0 IOPS 0 Bytes per-second`<br>`  Backward-rebuild-writes  0 IOPS 0 Bytes per-second`<br>`  Forward-rebuild-reads  0 IOPS 0 Bytes per-second`<br>`  Forward-rebuild-writes  0 IOPS 0 Bytes per-second`<br>`[root@Lnx-1 ~]# scli --query_protection_domain --protection_domain_name Domain1\|grep -i rebalance`<br>`  The number of parallel rebuild/rebalance jobs: 2`<br>`  Rebalance is Enabled and using Favor-Application-IO policy with the following parameters:`<br>`  The number of parallel rebuild/rebalance jobs: 2`<br>`  Rebalance is Enabled and using Favor-Application-IO policy with the following parameters:`<br>`  45.2 MB (46258 KB) active rebalance capacity`<br>`  1.6 GB (1608 MB) pending rebalance capacity`<br>`  1.6 GB (1653 MB) rebalance capacity`<br>`  Rebalance-reads  30 IOPS 30.8 MB (31539 KB) per-second`<br>`  Rebalance-writes  28 IOPS 28.4 MB (29081 KB) per-second`<br>`[root@Lnx-1 ~]#`<br><br>Besides the zero activity shown in the bottom lines of each command’s output, note the tuning parameters shown within the upper two lines of each output. |

| Step | Action |
|---|---|
| 10. | Return to (or re-open) the Remote Desktop session linked to Win-3.<br><br>Re-open the **Server Manager**, go to **Configuration/Services**, and find the two “emc” services used by ScaleIO. They are highlighted below:<br><br>Right-click them and select **Stop**.<br><br> 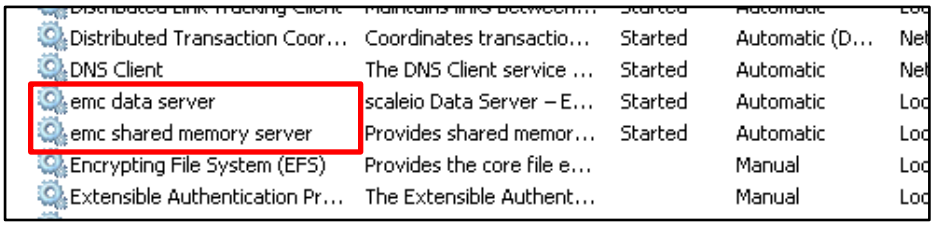 <br><br><br>The action we just performed will cause the SDS on Win3 to go offline. Note the **Alert** that shows up in the SDS pane of the GUI Dashboard view.<br><br> <br><br><br>You may also go to the GUI **Alert** view for more specific information.<br><br> 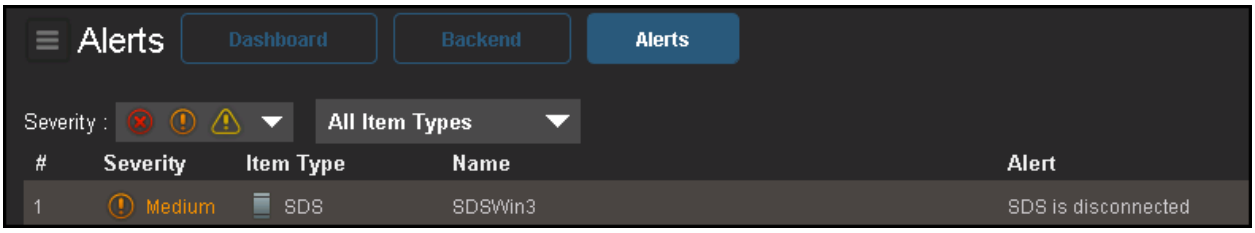 <br> |

| Step | Action |
|---|---|
| 11. | On the **Dashboard** view: note the Rebuild activity that is reported.<br> 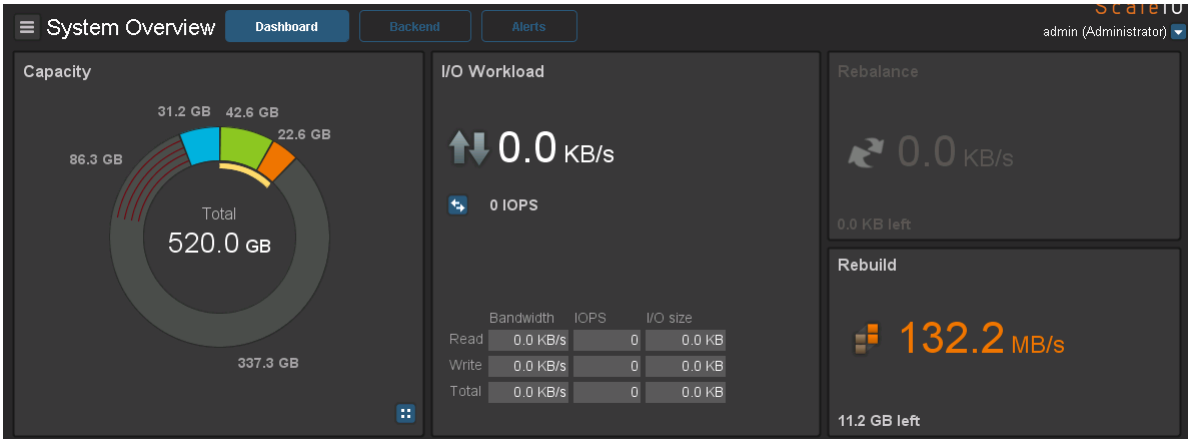 <br> |
| 12. | Go to the Lnx-1 Putty session and re-run the two commands showing Rebuild and Rebalance activity. What are the IOPS levels? Rebuild _____ Rebalance _____<br><br>[root@Lnx-1 ~]# scli --query_protection_domain --protection_domain_name Domain1\|grep -i rebuild<br>The number of parallel rebuild/rebalance jobs: 2<br>Rebuild is Enabled and using Limit-Concurrent-IO policy with the following parameters:<br>The number of parallel rebuild/rebalance jobs: 2<br>Rebuild is Enabled and using Limit-Concurrent-IO policy with the following parameters:<br>180.7 MB (184997 KB) active fwd-rebuild capacity<br>5.0 GB (5153 MB) pending fwd-rebuild capacity<br>0 Bytes active bck-rebuild capacity<br>0 Bytes pending bck-rebuild capacity<br>5.2 GB (5334 MB) fwd-rebuild capacity<br>0 Bytes bck-rebuild capacity<br>Backward-rebuild-reads 0 IOPS 0 Bytes per-second<br>Backward-rebuild-writes 0 IOPS 0 Bytes per-second<br>Forward-rebuild-reads 119 IOPS 119.8 MB (122675 KB) per-second<br>Forward-rebuild-writes 120 IOPS 120.2 MB (123084 KB) per-second<br>[root@Lnx-1 ~]# scli --query_protection_domain --protection_domain_name Domain1\|grep -i rebalance<br>The number of parallel rebuild/rebalance jobs: 2<br>Rebalance is Enabled and using Favor-Application-IO policy with the following parameters:<br>The number of parallel rebuild/rebalance jobs: 2<br>Rebalance is Enabled and using Favor-Application-IO policy with the following parameters:<br>0 Bytes active rebalance capacity<br>0 Bytes pending rebalance capacity<br>0 Bytes rebalance capacity<br>Rebalance-reads 0 IOPS 0 Bytes per-second<br>Rebalance-writes 0 IOPS 0 Bytes per-second<br>[root@Lnx-1 ~]#<br><br>**Note:** The **Backend** view of the Dashboard provides more detail regarding rebuild traffic. |

| Step | Action |
|---|---|
| 13. | Once the rebuilds complete, return to (or re-open) the Remote Desktop session to Win-3.<br><br>Re-open the **Server Manager**, go to **Configuration/Services**, and find the two “emc” services used by ScaleIO that you previously stopped.<br><br>Right-click them, and select **Start**.<br><br>Go to the dc Dashboard and note the rebuild and/or rebalance activity.<br><br> 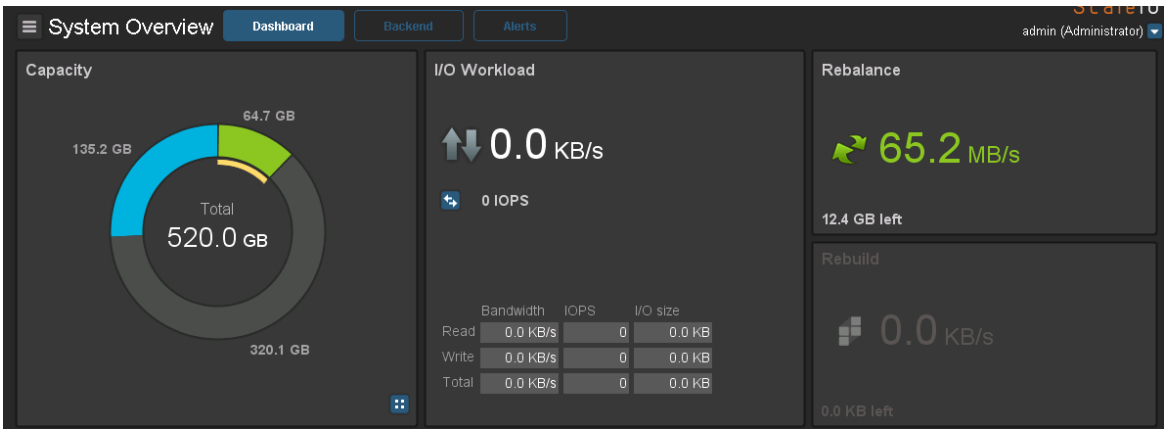 |
| 14. | Go to the Lnx-1 Putty session and re-run the two commands showing Rebuild and Rebalance activity. What are the IOPS levels? Rebuild _____ Rebalance _____<br><br>[root@Lnx-1 ~]# scli --query_protection_domain --protection_domain_name Domain1 --show_sds 1 \|grep -i rebuild<br>The number of parallel rebuild/rebalance jobs: 2<br>Rebuild is Enabled and using Limit-Concurrent-IO policy with the following parameters:<br>0 Bytes active fwd-rebuild capacity<br>0 Bytes pending fwd-rebuild capacity<br>0 Bytes active bck-rebuild capacity<br>0 Bytes pending bck-rebuild capacity<br>0 Bytes fwd-rebuild capacity<br>0 Bytes bck-rebuild capacity<br>Backward-rebuild-reads 0 IOPS 0 Bytes per-second<br>Backward-rebuild-writes 0 IOPS 0 Bytes per-second<br>Forward-rebuild-reads 0 IOPS 0 Bytes per-second<br>Forward-rebuild-writes 0 IOPS 0 Bytes per-second<br>[root@Lnx-1 ~]# scli --query_protection_domain --protection_domain_name Domain1 --show_sds 1 \|grep -i rebalance<br>The number of parallel rebuild/rebalance jobs: 2<br>Rebalance is Enabled and using Favor-Application-IO policy with the following parameters:<br>65.3 MB (66900 KB) active rebalance capacity<br>12.8 GB (13066 MB) pending rebalance capacity<br>12.8 GB (13131 MB) rebalance capacity<br>Rebalance-reads 20 IOPS 20.2 MB (20684 KB) per-second<br>Rebalance-writes 14 IOPS 14.6 MB (14950 KB) per-second |
| 15. | When the Rebalance completes, the task is complete. |

## Lab 4: Task 2 - Adjusting Rebuild and Rebalance Policy

| Step | Action |
|---|---|

| Step | Action |
|---|---|
| | You will use scli commands to change the rebuild and rebalance policies, then monitor the effects. |
| 1. | Once rebalance completes, adjust the rebuild policy to **no_limits**, and re-check the tuning parameters.<br><br>**`scli --set_rebuild_policy --protection_domain_name Domain1 --storage_pool_name Cap_Pool --policy no_limit`**<br><br>**`scli --query_protection_domain --protection_domain_name Domain1`**<br><br>`[root@Lnx-1 ~]# scli --set_rebuild_policy --protection_domain_name Domain1 --storage_pool_name Cap_Pool --policy no_limit`<br>`This setting might effect the speed of rebuild/rebalance, or degrade the performance of application IOs. Press 'y' to confirm that you want to proceed ?`<br>`Rebuild policy on Storage Pool Cap_Pool in Protection Domain Domain1 modified successfully`<br>`[root@Lnx-1 ~]# scli --query_protection_domain --protection_domain_name Domain1`<br>`Protection Domain Domain1 (Id: 3b45ef3100000000) has 1 storage pools, 1 Fault Sets, 5 SDS nodes, 4 volumes and 144.0 GB (147456 MB) available for volume allocation`<br>`Operational state is Active`<br><br>`Storage Pool Cap_Pool (Id: 1c09c63700000000) has 4 volumes and 144.0 GB (147456 MB) available for volume allocation`<br>`The number of parallel rebuild/rebalance jobs: 2`<br>`Rebuild is Enabled and using No-Limit policy`<br>`Rebalance is Enabled and using Favor-Application-IO policy with the following parameters:`<br>`Number of concurrent IOs per device: 1, Bandwidth limit per device: 10240 KB per second`<br>`Zero padding is disabled`<br>`Spare policy: 26% out of total`<br>`RAM cache is used`<br>`RAM cache write handling mode is 'passthrough'` |
| 2. | Return to (or re-open) the Remote Desktop session linked to Win-3.<br><br>Re-open the **Server Manager**, go to **Configuration/Services**, and find the two “emc” services used by ScaleIO. They are highlighted below:<br><br>Right-click them and select **Stop**.<br><br>Distributed Transaction Coor… Coordinates transactio… Started Automatic (D… Net<br>DNS Client The DNS Client service … Started Automatic Net<br>emc data server scaleio Data Server – E… Started Automatic Loc<br>emc shared memory server Provides shared memor… Started Automatic Loc<br>Encrypting File System (EFS) Provides the core file e… Manual Loc<br>Extensible Authentication Pr… The Extensible Authent… Manual Loc |

| Step | Action |
|---|---|
| 3. | Go to the dc Dashboard and note the rebuild and rebalance activity.<br> 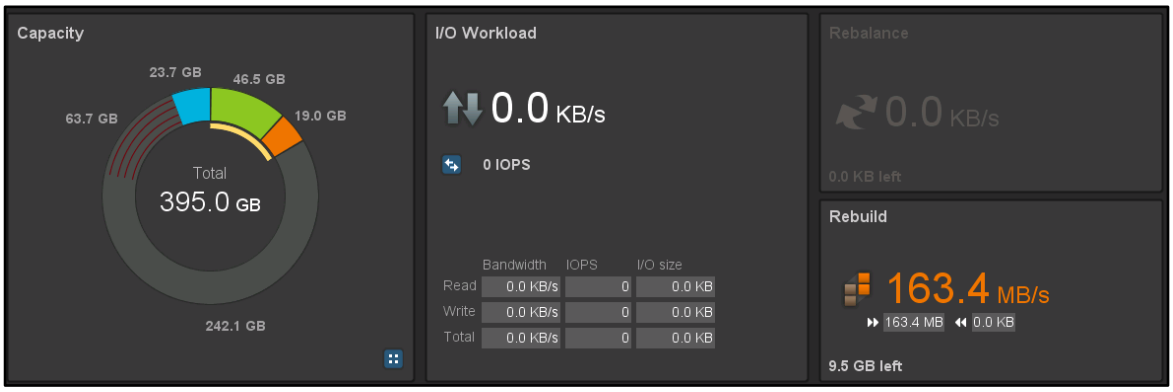 <br><br>What are the IOPS? Rebuild _____ Rebalance _____<br>Do you observe a change in IOPS relative to the constrained case in Task 1? _____ |
| 4. | Have a look at the dc GUI **Backend** view and observe the trend in the **Capacity In Use** column entry for each SDS server.<br>Are these trends what you would expect? _____ |
| 5. | Now, adjust the <u>rebalance</u> policy to **no_limits**, and re-check the tuning parameters.<br>**`scli --set_rebalance_policy --protection_domain_name Domain1 --storage_pool_name Cap_Pool --policy no_limit`**<br>**`scli --query_protection_domain --protection_domain_name Domain1 \|grep -i rebalance`**<br>`[root@Lnx-1 ~]# scli --set_rebalance_policy --protection_domain_name Domain1 --storage_pool_name Cap_Pool --policy no_limit`<br>`This setting might effect the speed of rebuild/rebalance, or degrade the performance of application IOs. Press 'y' to confirm proceed ?`<br>`Rebalance policy on Storage Pool Cap_Pool in Protection Domain Domain1 modified successfully`<br>`[root@Lnx-1 ~]# scli --query_protection_domain --protection_domain_name Domain1 \|grep -i rebalance`<br>`The number of parallel rebuild/rebalance jobs: 2`<br>`Rebalance is Enabled and using No-Limit policy`<br>`The number of parallel rebuild/rebalance jobs: 2`<br>`Rebalance is Enabled and using Favor-Application-IO policy with the following parameters:`<br>`0 Bytes active rebalance capacity`<br>`0 Bytes pending rebalance capacity`<br>`0 Bytes rebalance capacity`<br>`Rebalance-reads 0 IOPS 0 Bytes per-second`<br>`Rebalance-writes 0 IOPS 0 Bytes per-second` |
| 6. | Return to (or re-open) the Remote Desktop session linked to Win-3.<br>Re-open the **Server Manager**, go to **Configuration/Services**, and find the two “emc” services used by ScaleIO that you previously stopped.<br>Right-click them and select **Start**. |

| Step | Action |
|---|---|
| 7. | Go back to dc, open the ScaleIO GUI and note the rebuild and/or rebalance activity.<br><br>What are the IOPS? Rebuild _____ Rebalance _____<br><br>Did you get a greater IOPs rate for the Rebuild or the Rebalance? _____<br><br>Can you provide one possible reason for this? _____<br><br>(**Hint:** Make a list of the server-pairs involved during chunk copies in your case – first during the **Rebuild**, then during the **Rebalance**). |
| 8. | Choose the **Backend** view, by Storage Pool.<br><br>Right-click on the **Cap_Pool Storage Pool** and select **Set I/O Priority**.<br><br>Note that from this screen, you have additional tuning options for ScaleIO throttling.<br><br> 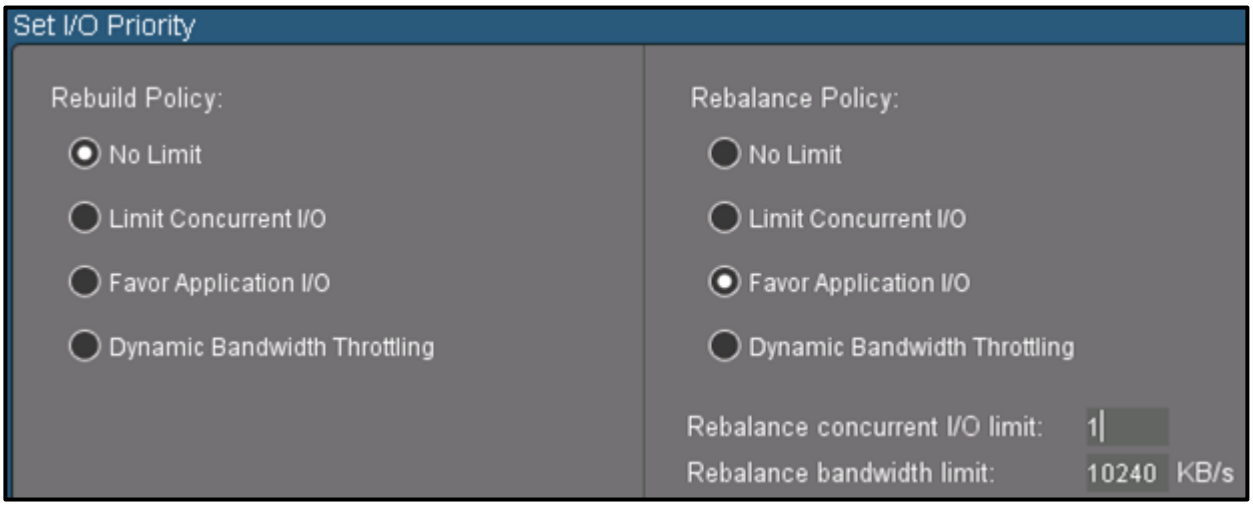 <br> |

## Lab 4: Task 3 - Configuring QoS for a Volume

| Step | Action |
|---|---|
| | Using scli, experiment with different volume limits. |

| Step | Action |
|---|---|
| 1. | From within a Putty session to the MDM Primary (should be Lnx-1), log into the MDM cluster as admin, if you are not admin already.<br><br>Run the following command:<br><br>**scli --query_sdc_volume_limits --volume_name vol2 --sdc_ip 192.168.1.11**<br><br>[root@Lnx-1 ~]# scli --query_sdc_volume_limits --volume_name vol2 --sdc_ip 192.168.1.11<br>sdc volume limits will be changed. Press 'y' to confirm that you want to proceed ?<br>Successfully queried limits from SDC 192.168.1.11 to volume vol2:IOPS limit is unlimited.Bandwidth limit is unlimited.<br><br>**Note:** The IP address 192.168.1.11 is Lnx-2, the server to which vol2 is currently mapped. |
| 2. | Return to (or re-open) a Putty session to Lnx-2.<br><br>On Lnx-2, generate some I/O activity to the ScaleIO volume. In your home directory **/root**, we have provided a script **writefile.bash** that will work fine for this purpose: it performs small-block sequential writes to a file called **largefile** in the filesystem layered on the ScaleIO device. When **largefile** reaches a filesize of ~4 GB, the script exits.<br><br>Examine, and then run, the script **writefile.bash** on Lnx-2:<br><br>**cd /root**<br><br>**ls –l writefile.bash**<br><br>**cat writefile.bash**<br><br>**./writefile.bash** |

| Step | Action |
|---|---|
| 3. | Look at the MB/sec or IOPS load shown on the Dashboard. *Remember that it's a running average, so allow the figures displayed to refresh two or three times for an accurate count.*<br> <br><br>**NOTE:** The **writefile.bash** script is designed to run only for a short period (minutes or less, depending on prevailing system load). If you notice the GUI is reporting 0 IOPs, that probably means the script has completed execution. As and when you need load on the ScaleIO volume **vol2**, rerun the **./writefile.bash** command. |
| 4. | Now, restrict the IOPS for the volume. Run this command on the Primary MDM:<br>**`scli --set_sdc_volume_limits --volume_name vol2 --sdc_ip 192.168.1.11 --limit_iops 10`**<br>`[root@Lnx-1 1]# scli --set_sdc_volume_limits --volume_name vol2 --sdc_ip 192.168.1.11 --limit_iops 10`<br>`Successfully modified limits from SDC 192.168.1.11 to volume vol2:IOPS limit set to 10.`<br>`[root@Lnx-1 1]#`<br>**Note:** An IOPS limit of 10 is the lowest possible setting. |
| 5. | Look at the IOPS load shown on the **Dashboard** view.<br>What is it now? _____<br>Experiment with higher IOPs settings – try 20 and 30, for example. Observe the changes.<br> 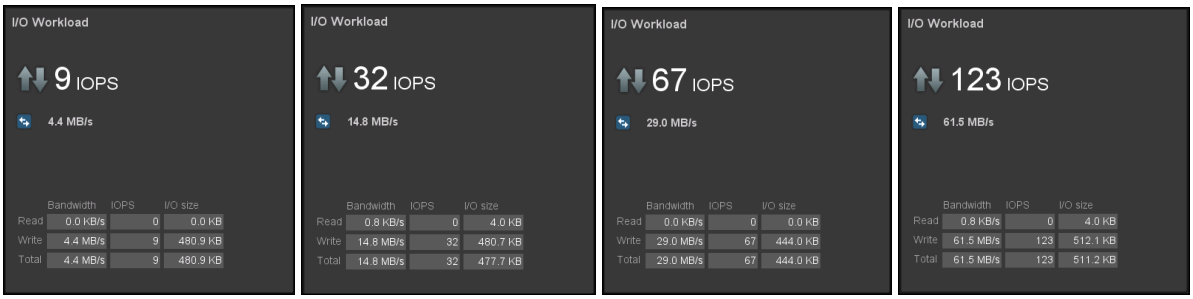 <br> |

| Step | Action |
|---|---|
| 6. | On Lnx-1: restore the volume configuration to the original state, before proceeding to the next task.<br><br>Return the IOPS limit to zero, which is unlimited.<br><br>`[root@Lnx-1 1]# scli --set_sdc_volume_limits --volume_name vol2 --sdc_ip 192.168.1.11 --limit_iops 0`<br>`Successfully modified limits from SDC 192.168.1.11 to volume vol2:IOPS limit set to unlimited.` |
| 7. | ***On Lnx-2: make sure to terminate your "writefile.bash" workload.*** This will help eliminate unnecessary load on the lab equipment.<br><br>On your PuTTY session to Lnx-2, repeatedly issue <**Ctrl>-C** until you get your **root@Lnx-2** shell prompt again. |

## Lab 4: Task 4 - RAM Cache Determination and Implementation

| Step | Action |
|---|---|
| | You will create a new storage pool, which will utilize RAM cache. |
| 1. | From within a Putty session to the MDM Primary (should be Lnx-1), log into the MDM cluster as admin, if you are not admin already.<br><br>Run the following command:<br><br>**`scli --query_protection_domain --protection_domain_name Domain1 \| less`** |
| 2. | Look at the information for the Storage Pool named **Cap_Pool**.<br><br>Is RAM cache being used for the Pool? Yes / No If so, what mode? _____<br><br>Look at the information for the SDS's.<br><br>Is the SDS RAM cache information listed? Yes / No |

| Step | Action |
|---|---|
| 3. | On the MDM, run the following command to look more closely at SDSWin2:<br><br>**`scli --query_sds --sds_name SDSWin2`**<br><br>Is the SDS RAM cache information listed?  Yes / No<br><br>If so,<br><br>1) What is its size in KB?  _____<br><br>2) Is it enabled?  Yes / No |
| 4. | Back on the dc server, in the Dashboard go to the **Backend** view.<br><br>Right-click on the **Domain1** and select **Add Storage Pool**.<br><br>The name of the new Storage Pool is **Perf_Pool**. Select **Use RAM Cache**.<br><br> 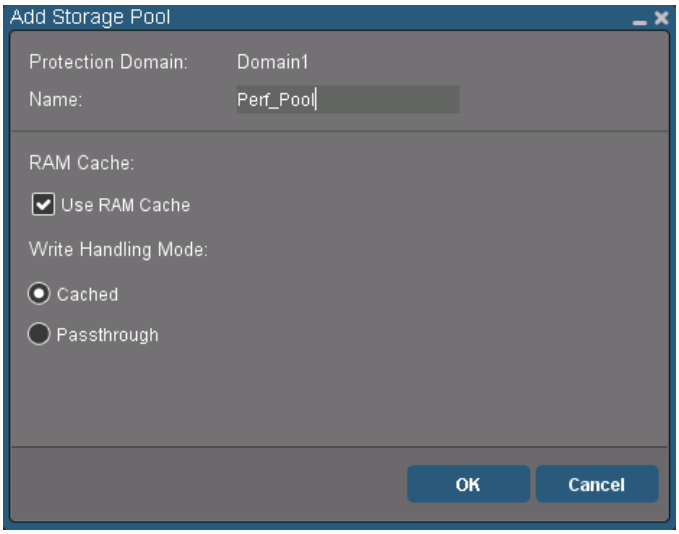 <br><br><br>When finished, click **OK**. Click **Close** when complete. |

| Step | Action |
|---|---|
| 5. | Back on the dc server, go to the ScaleIO GUI → **Backend** view. Sort by SDSs.<br><br>Right-click on the SDSLnx1 SDS and select **Add Device**.<br><br>The Device Path and Name is **/dev/sdc** and the Storage Pool is **Perf_Pool**.<br><br> <br><br><br>When finished, click **OK**. Click **Close** when complete.<br><br>Repeat these steps to add the **/dev/sdc** of SDSLnx2 and SDSLnx3 to **Perf_Pool**. |

| Step | Action |
|---|---|
| 6. | Check if any of your newly added devices has a **yellow** background (as below). If it does, it means that the device is not active, and should be **activated** (next step).<br> 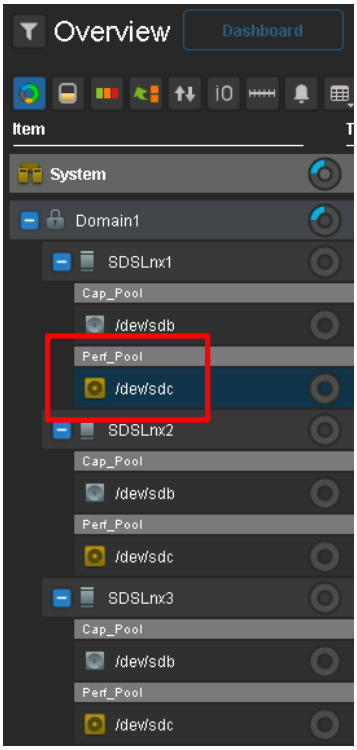 <br><br>If you don't see the yellow background, the devices are already active and ready for ScaleIO to use. Skip the next step and proceed directly to Step 7. |
| 7. | Activate all of the inactive devices - right-click on the device, and select **Activate Device**:<br> 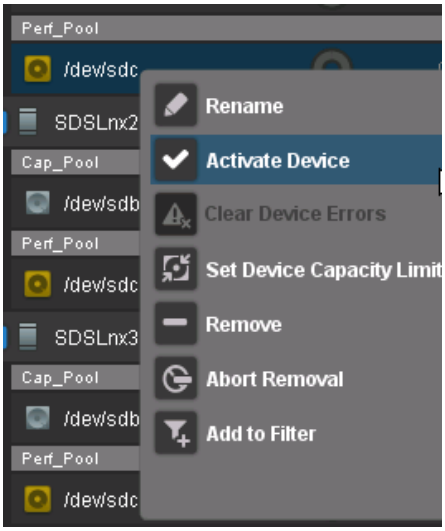 <br> |

| Step | Action |
|---|---|
| 8. | Back on the Primary MDM, run the following command:<br><br>**`scli --query_protection_domain --protection_domain_name Domain1 \| less`**<br><br>Look at the information for the Storage Pool named **Perf_Pool**.<br><br>Is RAM Cache being used for the Pool? Yes / No _____ |

## Lab 4: Task 5 - Monitoring ScaleIO-reported Events

| Step | Action |
|---|---|
| | |
| 1. | From within a Putty session to the MDM Primary (should be Lnx-1), log into the MDM cluster as admin, if you are not admin already.<br><br>**Note:** For a quick refresher on how to do this, refer to Lab3, Tasks 2 and 3.<br><br>Run the following command:<br><br>**`cd /opt/emc/scaleio/mdm/logs`**<br><br>**`ls -al`**<br><br>`[root@Lnx-1 logs]# cd /opt/emc/scaleio/mdm/logs`<br>`[root@Lnx-1 logs]# ls -al`<br>`total 105656`<br>`drwx------. 2 root root  4096 Jun 24 14:20 .`<br>`drwxr-xr-x. 7 root root  4096 Jun 19 11:35 ..`<br>`-rw-r--r--. 1 root root  199 Jun 19 11:35 emc-scaleio-mdm-install.log`<br>`-rw-r--r--. 1 root root  25600 Jun 24 14:20 eventlog.db`<br>`-rw-r--r--. 1 root root  289646 Jun 24 14:23 exp.0`<br>`-rw-r--r--. 1 root root  1200070 Jul  1 08:53 trc.0`<br>`-rw-r--r--. 1 root root 26217908 Jun 24 13:25 trc.1`<br>`-rw-r--r--. 1 root root 27548551 Jun 24 13:13 trc.2`<br>`-rw-r--r--. 1 root root 26292469 Jun 24 13:11 trc.3`<br>`-rw-r--r--. 1 root root 26583384 Jun 24 13:09 trc.4` |

| Step | Action |
|---|---|
| 2. | This is where ScaleIO keeps its logs. ScaleIO sees certain activities (events) as worthy of note, and notes them within **eventlog.db**.<br><br>Still in the same directory, run the following command:<br><br>**`file eventlog.db`**<br><br>What type of file is **eventlog.db**? _____ |
| 3. | Still in the same directory, run the following command:<br><br>**`more eventlog.db`**<br><br>While you can read some of this, it is not in a helpful format. Fortunately, ScaleIO has a tool to assist with this.<br><br>Now, switch directories and run:<br><br>**`cd /opt/emc/scaleio/mdm/bin`**<br><br>**`./showevents.py`**<br><br>Some of the events displayed should look familiar. |

| Step | Action |
|---|---|
| 4. | Now run the same command with –help:<br><br>**./showevents.py --help**<br><br>[root@Lnx-1 bin]# ./showevents.py --help<br>Usage: showevents.py [options]<br><br>Options:<br>-h, --help show this help message and exit<br>--min_severity=MIN_SEVERITY<br>Specify minimal severity<br>--severity=SEVERITY Specify severity<br>--from_id=ID_FROM Limit by ID<br>--to_id=ID_TO Limit by ID<br>--from_date=DATE_FROM<br>Limit by from Date. Date format is same as it shown in event log (YYYY-MM-DD HH:MM:SS) For example January 1st 2013 09:00PM is 2013-01-01 09:00:00<br>--to_date=DATE_TO Limit by to Date. Date format is same as it shown in event log (YYYY-MM-DD HH:MM:SS) For example January 1st 2013 09:00PM is 2013-01-01 09:00:00<br>-g GREP, --grep=GREP Search specific word<br>--full Show full event<br><br>As you can see, you can filter events by date range, by severity, and even search (**grep**) for key words. |
| 5. | Return to (or re-open) the Remote Desktop session linked to Win-3.<br><br>Re-open the **Server Manager**, go to **Configuration/Services**, and find the two "emc" services used by ScaleIO. They are highlighted below:<br><br>Right-click them and select "Stop".<br><br>Distributed Transaction Coor… Coordinates transactio… Started Automatic (D… Net<br>DNS Client The DNS Client service … Started Automatic Net<br>emc data server scaleio Data Server – E… Started Automatic Loc<br>emc shared memory server Provides shared memor… Started Automatic Loc<br>Encrypting File System (EFS) Provides the core file e… Manual Loc<br>Extensible Authentication Pr… The Extensible Authent… Manual Loc |
| 6. | Go to the Lnx-1 Putty session and run the command to view the log:<br><br>**./showevents.py**<br><br>173 2014-07-02 11:16:51.349 CLI_COMMAND_SUCCEEDED INFO Command modify_sdc_volume_limit succeeded.<br>174 2014-07-02 11:17:05.300 MDM_CLI_COMMAND_RECEIVED INFO Command modify_sdc_volume_limit received, User: 'admin'.<br>175 2014-07-02 11:17:05.386 CLI_COMMAND_SUCCEEDED INFO Command modify_sdc_volume_limit succeeded.<br>176 2014-07-02 11:30:48.147 SDS_DISCONNECTED ERROR SDS: SDSWin3 IP: 192.168.1.22 (id: 8c79d30600000005) disconnected.<br>177 2014-07-02 11:30:48.998 MDM_DATA_DEGRADED ERROR The system is now in DEGRADED state. |

| Step | Action |
|---|---|
| 7. | Return to the Win-3 Remote Desktop session and start the emc services.<br><br>Go to the Lnx-1 Putty session and run the command to view the log:<br><br>**./showevents.py**<br><br>176 2014-07-02 11:30:48.147 SDS_DISCONNECTED ERROR SDS: SDSWin3 IP: 192.168.1.22 (id: 8c79d30600000005) disconnected.<br>177 2014-07-02 11:30:48.998 MDM_DATA_DEGRADED ERROR The system is now in DEGRADED state.<br>178 2014-07-02 11:33:29.723 MDM_DATA_NORMAL INFO The system is now in NORMAL state.<br>179 2014-07-02 11:34:21.909 SDS_RECONNECTED INFO SDS: SDSWin3 (ID 8c79d30600000005) reconnected |

## Lab 4: Task 6 - Set Up Centralized RSyslog Server for ScaleIO Event Logging

| Step | Action |
|---|---|
| | |
| 1. | Our first step is to identify the host that we will configure as our centralized **RSYSLOG** server. This is the target syslog server to which ScaleIO events will be redirected.<br><br>For this lab we will use the ScaleIO **TB** server (Lnx-3, 192.168.1.12) as our target server.<br><br>Using PuTTY, login to the Lnx-3 server with the credentials **root/P@ssw0rd**. |
| 2. | On Lnx-3, check the status of the Linux firewall – the status right now, and the expected status after a reboot:<br><br>**service iptables status**<br><br>Observe that the firewall is currently not running.<br><br>**chkconfig --list iptables**<br><br>The output of chkconfig ("off" for all runlevels 0 through 6) tells us that the firewall service will never automatically start on this host, whenever it starts up.<br><br>This means that Lnx-3 has unrestricted access on both incoming and outgoing TCP traffic. **<u>You can proceed to Step 3 right now.</u>** |
| 3. | On Lnx-3, examine the **/etc/rsyslog.conf** file, and make suitable changes to it, following the instructions provided below.<br><br>Take a look at the default configuration file **/etc/rsyslog.conf** :<br><br>**more /etc/rsyslog.conf**<br><br>and page through to the end of the file, using the **<Space>** bar. |

| Step | Action |
|---|---|
| 4. | For your convenience, we have provided you “before” and “after” versions of the **rsyslog.conf** file to enable centralized logging for ScaleIO. These versions are in the **/root** home directory for your reference:<br><br>**`ls -l /root/rsyslog*`**<br><br>You should see two rsyslog-related files listed in root’s home directory - **rsyslog.conf.orig** and **rsyslog.conf.new**, as shown below:<br><br>`[root@Lnx-3 etc]# ls -l /root/rsyslog*`<br>`-rw-r--r--. 1 root root 2826 Aug 18 16:05 /root/rsyslog.conf.new`<br>`-rw-r--r--. 1 root root 2741 May 20  2009 /root/rsyslog.conf.orig`<br>`[root@Lnx-3 etc]#` |
| 5. | Ensure that the file **/root/rsyslog.conf.orig** is exactly the same as your current **/etc/rsyslog.conf**:<br><br>**`diff /etc/rsyslog.conf /root/rsyslog.conf.orig`**<br><br>No differences should be reported. |
| 6. | On Lnx-3, examine the needed changes in the **rsyslog.conf** file:<br><br>**`diff /root/rsyslog.conf.orig /root/rsyslog.conf.new`**<br><br>`[root@Lnx-3 ~]# diff /root/rsyslog.conf.orig /root/rsyslog.conf.new`<br>`13,14c13,14`<br>`< #$ModLoad imudp.so`<br>`< #$UDPServerRun 514`<br>`---`<br>`> $ModLoad imudp.so`<br>`> $UDPServerRun 514`<br>`17,18c17,18`<br>`< #$ModLoad imtcp.so`<br>`< #$InputTCPServerRun 514`<br>`---`<br>`> $ModLoad imtcp.so`<br>`> $InputTCPServerRun 514`<br>`39c39`<br>`< *.info;mail.none;authpriv.none;cron.none  /var/log/messages`<br>`---`<br>`> *.info;mail.none;authpriv.none;cron.none;local0.none  /var/log/`<br>`messages`<br>`60c60,61`<br>`<`<br>`---`<br>`> $template ScaleIOMsg, "/var/log/%PROGRAMNAME%.log"`<br>`> local0.*  ?ScaleIOMsg` |

| Step | Action |
|---|---|
| 7. | Observe the following differences.<br><br>The first four modified lines (beginning with **$ModLoad** or **$Input)** are just uncommented lines, i.e., the preceding **#** has been removed.<br><br>The redirection to the **/var/log/messages** file has been enhanced to include the **local0** facility – this is the rsyslog facility we will use for ScaleIO event logging.<br><br>The last two lines for the ScaleIO rsyslog template are new. They indicate that facility **local0** should be used for ScaleIO logging, and specify that the logging should be directed to a file in the **/var/log** directory with the file name as provided. |
| 8. | On Lnx-3, now that we understand the needed changes, let's actually make them and activate them.<br><br>First, replace the **/etc/rsyslog.conf** file with the modified version:<br><br>**`cp /root/rsyslog.conf.new /etc/rsyslog.conf`**<br><br>Next, restart the **RSYSLOG** service for the configuration changes to take effect:<br><br>**`service rsyslog restart`** |
| 9. | On the target RSYSLOG server Lnx-3, check if the server is listening on port 514:<br><br>**`netstat -uanp \| grep rsyslog`**<br><br>You should see results as below:<br><br>**`udp  0  0 0.0.0.0:514  0.0.0.0:*  3589/rsyslogd`**<br>**`udp  0  0 :::514  :::*  3589/rsyslogd`**<br><br>This completes the configuration for the centralized target logging server. Next, let's configure the ScaleIO cluster to direct event logging traffic to the target server. |
| 10. | On the Primary MDM server Lnx-1, run the following **scli** commands, to perform the desired redirection for rsyslog (making sure to first login as **admin/P@ssw0rd** to ScaleIO):<br><br>**`scli --login --username admin --password P@ssw0rd`**<br><br>**`scli --stop_remote_syslog`**<br><br>**`scli --start_remote_syslog --remote_syslog_server_ip 192.168.1.12 --remote_syslog_server_port  514`** |

| Step | Action |
|---|---|
| 11. | On primary MDM server Lnx-1: perform similar steps as we did previously for the target host Lnx-3 (steps 2 through 5, listed below for your convenience):<br><br>**`service iptables status`**<br><br>**`chkconfig --list iptables`**<br><br>The above commands confirm that the firewall will not block any outgoing traffic from Lnx-1.<br><br>**`diff /etc/rsyslog.conf /root/rsyslog.conf.orig`**<br><br>This shows the **rsyslog.conf.orig** file is the same as the currently active configuration file.<br><br>**`diff /root/rsyslog.conf.orig /root/rsyslog.conf.new`**<br><br>This will show only the following change, which clearly suggests that all syslog traffic from this host, Lnx-1, is to be redirected to the Lnx-3 host on TCP port 514 (which is the port it is listening on):<br><br>**`< #*.* @@remote-host:514`**<br>**`> *.* @@192.168.1.12:514`**<br><br>**`cp /root/rsyslog.conf.new /etc/rsyslog.conf`**<br><br>**`service rsyslog restart`**<br><br>The two commands above will cause, and activate, the needed change to **rsyslog.conf** on the Lnx-1 host. |
| 12. | Repeat all commands within Step 8 on the secondary MDM server Lnx-2.<br><br>What is the purpose of doing this? Wouldn't the primary server Lnx-1 generate all the events to be logged? _____ |

| Step | Action |
|---|---|
| 13. | Now, let's validate that the centralized rsyslog server Lnx-3 is receiving logs from the clients.<br><br>First, confirm that Lnx-1 is the primary MDM, and Lnx-2 is the secondary MDM. You can do this with the **`scli --query_cluster`** command. The Primary MDM will result in output similar to this:<br><br>`[root@Lnx-1 ~]# scli --query_cluster`<br>` Mode: Cluster, Cluster State: Normal, Tie-Breaker State: Connected`<br>` Primary IP: 192.168.1.10`<br>` Secondary IP: 192.168.1.11`<br>` Tie-Breaker IP: 192.168.1.12`<br>` Management IP: 192.168.1.10, 192.168.1.11`<br>` Name: N/A`<br>`[root@Lnx-1 ~]#`<br><br>If you run the same **`scli --query_cluste`**r command on the Secondary MDM, you will see output similar to this:<br><br>`[root@Lnx-2 ~]# scli --query_cluster`<br>`Error: MDM failed command.  Status: Use the cluster virtual IP or the primary MDM`<br>` physical IP to access the cluster` |
| 14. | On the Primary MDM (Lnx-1, 192.168.1.10), run the command:<br><br>**`scli --switch_mdm_ownership`** |
| 15. | Wait about 15 seconds and check the **/var/log** folder on the central rsyslog server (Lnx-3 in this exercise). You should see a file named **EMC-ScaleIO.log**. |
| 16. | On Lnx-3, check the contents of this log file. You should see entries representing the MDM switch ownership commands just run. In the example below, Lnx-2 (192.168.1.11) is now the Primary MDM.<br><br>`Aug  5 04:20:23 Lnx-2 EMC-ScaleIO -[ Initialized the remote syslog module ]`<br>`Aug  5 12:14:28 Lnx-2 EMC-ScaleIO -[ This MDM is switching to Master mode. MDM will start running. ]`<br>`Aug  5 12:14:33 Lnx-2 EMC-ScaleIO -[ SDS: SDS_Lnx2 (ID 95c0773b00000000) reconnected ]`<br>`Aug  5 12:14:33 Lnx-2 EMC-ScaleIO -[ SDS: SDS_Lnx1 (ID 95c0773d00000002) reconnected ]`<br>`Aug  5 12:14:33 Lnx-2 EMC-ScaleIO -[ SDS: SDS_Lnx3 (ID 95c0773c00000001) reconnected ]`<br>`Aug  5 12:14:33 Lnx-2 EMC-ScaleIO -[ SDS: SDS_Win2 (ID 95c09e4b00000004) reconnected ]`<br>`Aug  5 12:14:33 Lnx-2 EMC-ScaleIO -[ SDC connected. ID: 64022b1a00000001; IP: 192.168.1.11; GUID: 1DDB2C2E-2D02-41AE`<br>`-BB28-10EEF0537A0C ]`<br>`Aug  5 12:14:33 Lnx-2 EMC-ScaleIO -[ SDC connected. ID: 64022b1b00000002; IP: 192.168.1.12; GUID: BD9DAA80-16C3-4043`<br>`-AB11-380F083A22DE ]`<br>`Aug  5 12:14:33 Lnx-2 EMC-ScaleIO -[ SDC connected. ID: 64022b1900000000; IP: 192.168.1.10; GUID: FE5CA99B-4D67-46D4`<br>`-A09C-A8C6C36AADFB ]` |

| Step | Action |
|---|---|
| 17. | On the new primary MDM Lnx-2: run the command<br><br>**`scli --switch_mdm_ownership`**<br><br>Note that the centralized log file **/var/log/EMC-ScaleIO.log** shows that the Primary MDM is now back to Lnx-1 (192.168.1.10):<br><br>`Aug  5 12:14:28 Lnx-1 EMC-ScaleIO -[ Initialized the remote syslog module ]`<br>`Aug  5 12:19:09 Lnx-1 EMC-ScaleIO -[ This MDM is switching to Master mode. MDM will start running. ]`<br>`Aug  5 12:19:12 Lnx-1 EMC-ScaleIO -[ SDS: SDS_Lnx2 (ID 95c0773b00000000) reconnected ]`<br>`Aug  5 12:19:12 Lnx-1 EMC-ScaleIO -[ SDS: SDS_Lnx3 (ID 95c0773c00000001) reconnected ]`<br>`Aug  5 12:19:12 Lnx-1 EMC-ScaleIO -[ SDS: SDS_Lnx1 (ID 95c0773d00000002) reconnected ]`<br>`Aug  5 12:19:12 Lnx-1 EMC-ScaleIO -[ SDS: SDS_Win2 (ID 95c09e4b00000004) reconnected ]`<br>`Aug  5 12:19:12 Lnx-1 EMC-ScaleIO -[ SDC connected. ID: 64022b1900000000; IP: 192.168.1.10; GUID: FE5CA99B-4D67-46`<br>`-A09C-A8C6C36AADFB ]`<br>`Aug  5 12:19:12 Lnx-1 EMC-ScaleIO -[ SDC connected. ID: 64022b1b00000002; IP: 192.168.1.12; GUID: BD9DAA80-16C3-40`<br>`-AB11-380F083A22DE ]`<br>`Aug  5 12:19:12 Lnx-1 EMC-ScaleIO -[ SDC connected. ID: 64022b1a00000001; IP: 192.168.1.11; GUID: 1DDB2C2E-2D02-41`<br>`-BB28-10EEF0537A0C ]`<br>`Aug  5 12:19:13 Lnx-1 EMC-ScaleIO -[ SDC connected. ID: 6402522900000003; IP: 192.168.1.21; GUID: 62F30EE7-E8B3-C7`<br>`-B1A0-1FC716FF8EF3 ]`<br><br>This confirms that both Lnx-1 and Lnx-2 are successfully performing redirection of ScaleIO events whenever either one becomes the primary. |

## Lab Exercise 5-A: Exercising the ScaleIO REST API

| | |
|---|---|
| **Purpose:** | Use a graphical browser-based REST client to issue standard REST calls to a ScaleIO cluster and verify the results. |

| | |
|---|---|
| **Tasks:** | ✓ Validate that the REST gateway has a functional connection to the ScaleIO MDM<br>✓ Install and prepare a Firefox browser-based REST client<br>✓ Issue "GET" and "POST" calls to a ScaleIO instance using the REST client |

| | |
|---|---|
| **Reference:** | ScaleIO Lab Diagram (graphically showing relevant details for all VMs within a student pod – refer to Page 2) |

## Lab 5-A: Task 1 - Verify REST Gateway Connection to the ScaleIO Instance

| Step | Action |
|---|---|
| | |
| 1. | Open a Browser on the DC. Connect to https://dc.emc.edu/rest.jsp |
| 2. | It will prompt you to logon to the Installer: use U: admin \| P: P@ssw0rd 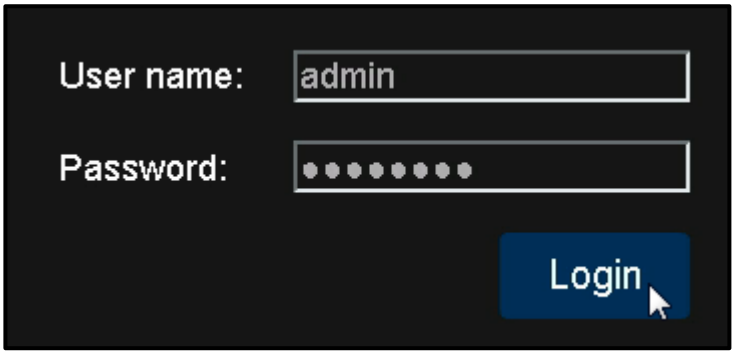  |
| 3. | It will prompt you to connect to the REST Gateway. Click the link 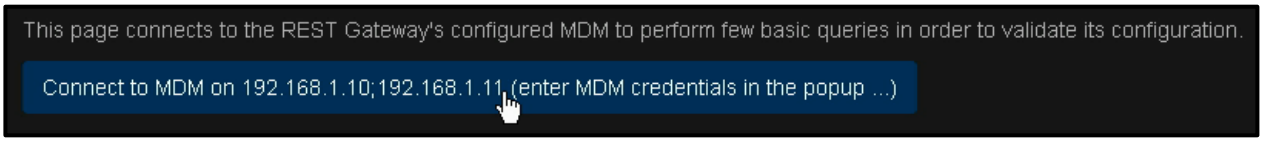  |
| 4. | Enter admin \| P@ssw0rd as the credentials (these are the MDM credentials) 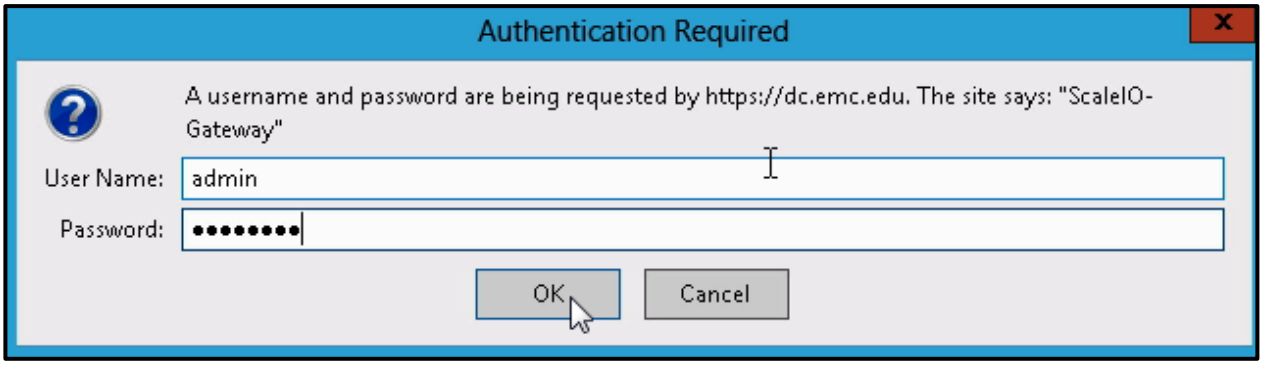  |

| Step | Action |
|---|---|
| 5. | It will automatically run some basic queries, and validate the REST configuration<br> 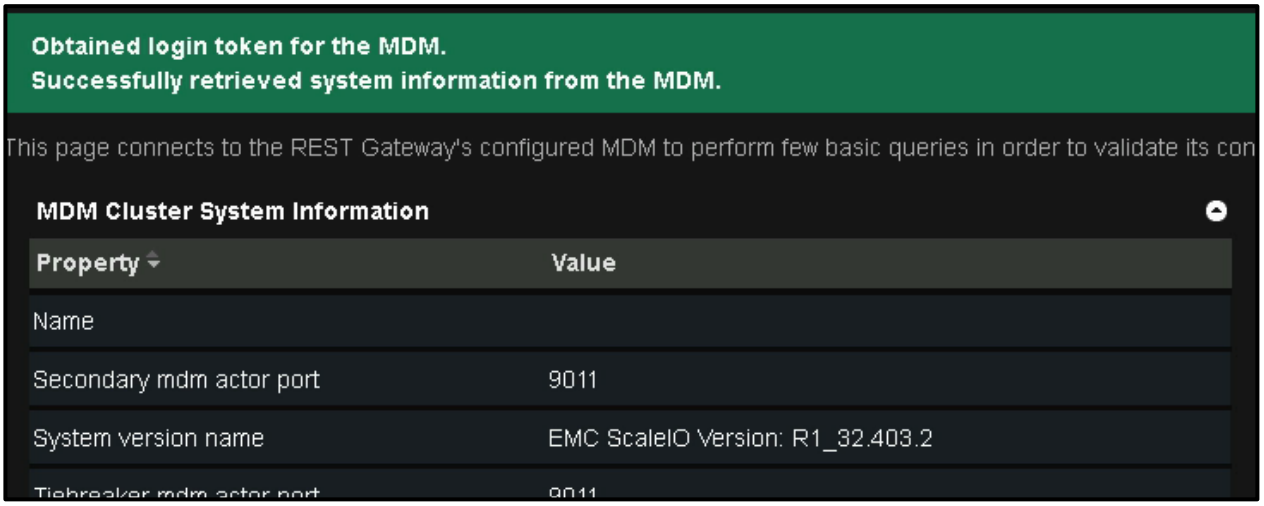 <br> |

## Lab 5-A: Task 2 - Prepare the Firefox REST Client

| Step | Action |
|---|---|
| | |
| 1. | In the Firefox browser, click on the red RESTClient link, in the upper-right corner<br> 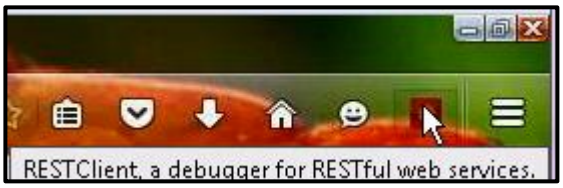 <br> |
| 2. | It will launch the client<br> <br> |

| Step | Action |
|---|---|
| 3. | We've prepared some requests in advance, click Favorite Requests → Import Favorite Requests 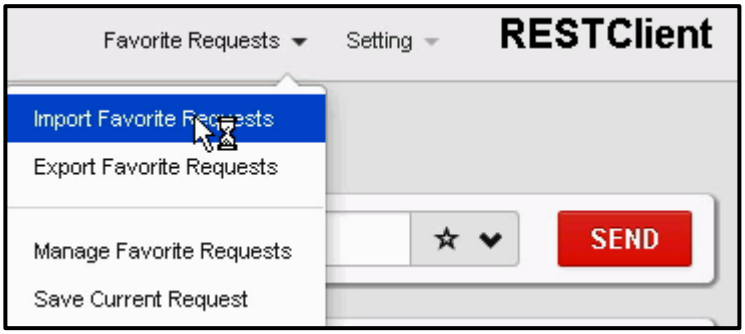  |
| 4. | Browse to c:\share\misc, select" Validated…json" and click Open 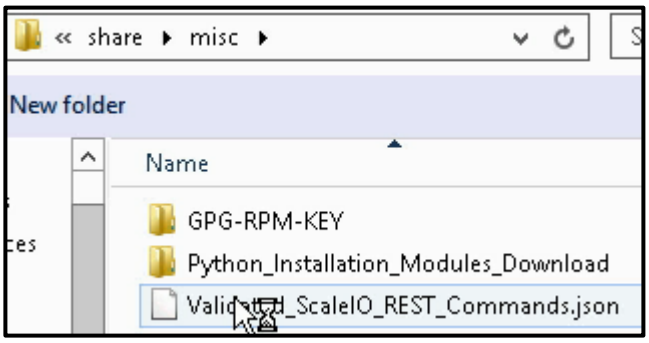  |
| 5. | They will not appear (minor bug in client). Repeat step 4, and they will appear 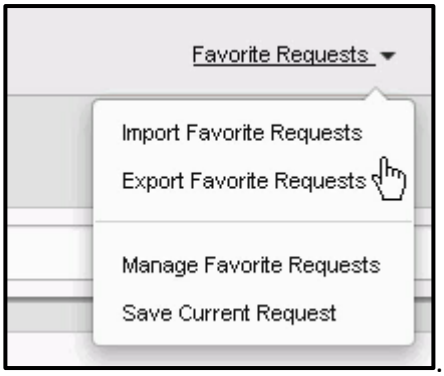  … |

## Lab 5-A: Task 3 - Issue “GET” and “POST” Requests to ScaleIO

| Step | Action |
|---|---|
| | |
| 1. | Load our first practice REST call, **get_gateway_configuration**.<br><br>From **Favorite Requests**, select the preloaded request **get_gateway_configuration**.<br><br>Observe the three items:<br> 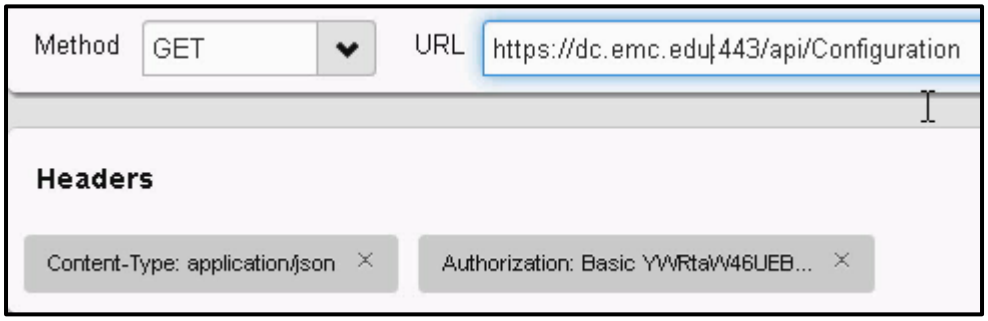 <br><br><br>Note that the REST Method is automatically selected as **GET** for this saved call. Note the syntax of the URL to get the REST gateway’s configuration.<br><br>Also, note that the **Content-Type** shows the value **application-json.** If you ever need to change to another type, you can use the **Headers** drop-down for doing this. |
| 2. | Specify the credentials for your REST gateway.<br><br>The saved query that we just imported contains gateway credentials from another installation. Since we need to authenticate to our specific REST gateway, we need to change to our gateway’s credentials. Do this as follows:<br><br>First, use the **X** button shown, to clear out the existing credentials:<br> <br> |

| Step | Action |
|---|---|
| 3. | Next, click on **Authentication** (on the Menu bar above) and select **Basic Authentication**. In the next screen offered, specify the REST gateway username **admin** and the password **P@ssw0rd**. Check the **Remember Me** box and click **OK**.<br> 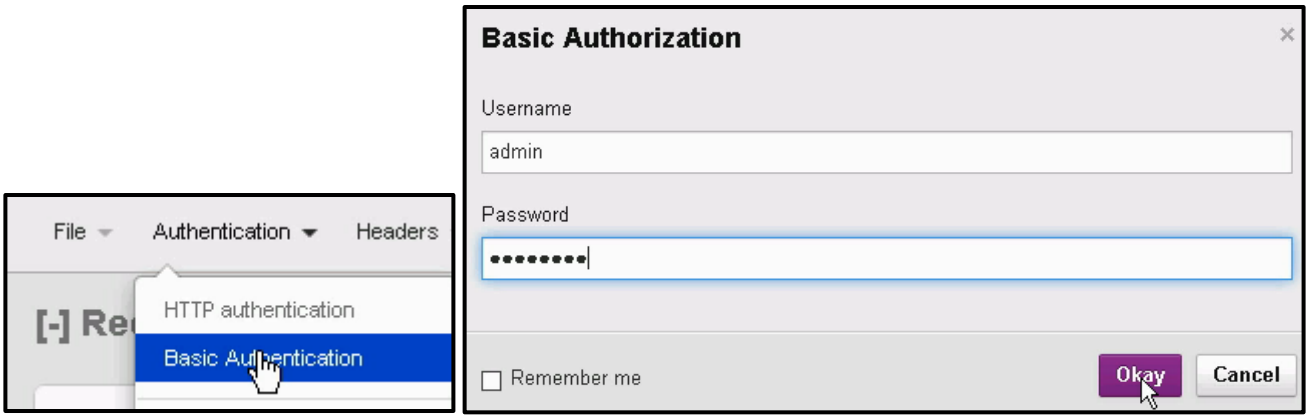 <br><br>**Note:** You may need to correct the credentials in this manner, before running ***every*** other imported query in your Favorite Requests, as you proceed through this lab exercise. |
| 4. | Run the **get_gateway_configuration** query.<br><br>Make sure the BODY section is left blank. This is true for all GET requests.<br><br>Now that the query is filled out, click **SEND**, to the right of the URL. |

| Step | Action |
|---|---|
| 5. | Examine the response: check that the REST gateway is configured with the correct IP addresses for the MDM.<br><br>Look at the **Response** section of the REST client, below **Body**.<br><br>Select the **Response Body (Highlight)** section, as shown here:<br><br>Response Headers \| Response Body (Raw) \| Response Body (Highlight) \| Response Body (Preview)<br><br>1. {<br>2. "snmpResendFrequency": "0",<br>3. "systemId": "4b36f8677d4ecb5d",<br>4. "snmpSamplingFrequency": "30",<br>5. "mdmPort": "6611",<br>6. "featuresEnableSnmp": "false",<br>7. "featuresEnableIM": "true",<br>8. "snmpPort": "162",<br>9. "mdmUsername": "",<br>10. "mdmAddresses":<br>11. [<br>12. "192.168.1.10",<br>13. "192.168.1.11"<br>14. ],<br>15. "snmpTrapsReceiverIps": null<br><br>In the response, ensure that your REST gateway correctly reports the primary and secondary MDM addresses as shown. This confirms that the gateway is properly configured for our specific MDM instance. |
| 6. | Select **get_all_instances** from **Favorite Requests,** and run this new request.<br><br>Method GET URL https://dc.emc.edu:443/api/types/System/instances<br><br>**Headers**<br><br>Content-Type: application/json × Authorization: Basic YWRtaW46UEB... × |

| Step | Action |
|---|---|
| 7. | Examine the information reported in the **Response Body (Highlight)** section to answer these questions:<br><br>What is the installation ID of the reported instance? _____<br><br>Which is the version number of this instance? _____<br><br>What are the IP addresses of the three MDM cluster nodes? _____<br><br>_____ |
| 8. | **Note:** Make sure to record the **id** from the output as indicated here. Also, highlight it and copy/paste to a Notepad window, as it will be needed later (in a POST request):<br>"qosTimeLeftInDays": "5",<br>"replicationTimeLeftInDays": "5",<br>"isInitialLicense": true,<br>"id": "7843f2302198b8b8",<br>"links":<br>[<br>{<br>"rel": "self",<br>"href": "/api/instances/System::7843f2302198b8b8"<br>},<br><br>What is the **id** of your ScaleIO instance? _____ |
| 9. | Get the Protection Domains configured for this instance.<br><br>Load the **get_protdoms_this_instance** command from **Favorite Requests**. Examine the URL and then run it by clicking **SEND**.<br><br>URL https://dc.emc.edu:443/api/types/ProtectionDomain/instances |
| 10. | Review the results reported for the Protection Domain(s) in the **Response Body (Highlight)** pane.<br><br>How many Protection Domains are reported? _____<br><br>Are the name(s) of the Protection Domain(s) reported? _____<br><br>Are you sure that the Protection Domain corresponds to your particular ScaleIO installation, and not someone else's? Why? _____ |

| Step | Action |
|---|---|
| 11. | Get the volumes that have been created in this instance.<br><br>Next, load and run the **get_volumes_this_instance** query from **Favorite Requests**.<br><br>URL https://dc.emc.edu:443/api/types/Volume/instances |
| 12. | In the **Response Body (Highlight)** pane, locate an easily identifiable volume with a name that you assigned and you can recognize (for example, **vol4** below):<br><br>{<br>"name": "vol4",<br>"creationTime": 1409256020,<br>"mappingToAllSdcsEnabled": false,<br>"consistencyGroupId": "0000000000000000",<br>"ancestorVolumeId": "0000000000000000",<br>"mappedScsiInitiatorInfo": null,<br>"storagePoolId": "ce32103700000000",<br>"isObfuscated": false,<br>"volumeType": "ThickProvisioned",<br>"vtreeId": "3e4e4d0400000003",<br>"sizeInKb": 16777216,<br>"mappedSdcInfo": null,<br>"id": "f663ed5600000003",<br>"links":<br><br>Using the mouse, highlight the **id** reported for this volume (hex string within the quotation marks). Right-click and select **Copy**, then **Paste** to the Notepad buffer where the instance ID was saved. The Volume ID will also be needed in our POST request.<br><br>We have exercised a few basic GET calls to the instance. Let's try our POST call next. |

| Step | Action |
|---|---|
| 13. | In the Firefox REST Client, load **create_snapshot** from **Favorite Requests**.<br><br>Check first that the command type has now changed to **POST** (it is not a **GET**).<br><br>Using the saved results (in Notepad window), make the needed edits in the URL and in the BODY section of the POST request.<br><br>In the screen shown next, there are ***three*** items you will need to change: the **instance ID** in the URL; **volume ID**, and a **name** you assign for the new snapshot:<br><br>Method POST URL https://192.168.1.20:443/api/instances/System:452a6a6a45371112/action/snapshotVolumes<br><br>**Headers**<br>Content-Type: application/json<br>Authorization: Basic YWRtaW46UEB...<br><br>**Body**<br>{<br>"snapshotDefs":[<br>{"volumeId":"f663ed5600000003","snapshotName":"SNAP_vol4"}<br>]<br>}<br><br>Click **SEND** to run the REST command. |
| 14. | Check that no errors are reported. Correct the query and rerun, if needed.<br><br>Examine the contents of the **Response Body (Highlight)** pane. You should not see any errors being reported.<br><br>If the command ran successfully, this is what the results will look like:<br><br>[-] Response<br>Response Headers \| Response Body (Raw) \| Response Body (Highlight) \| Response Body (Preview)<br><br>1. {<br>2. "snapshotGroupId": "edaa663c00000001",<br>3. "volumeIdList":<br>4. [<br>5. "b85dcf2700000001"<br>6. ]<br>7. } |

| Step | Action |
|---|---|
| 15. | Using an SCLI command, verify that the **REST POST** request has run successfully.<br><br>On an ssh session, as root to your Primary MDM (either **192.168.1.10** or **192.168.1.11**), first authenticate as user **admin**; then run the command:<br><br>**`scli --query_all_volumes`**<br><br>Can you identify the new snapshot within the list of volumes that are reported?<br><br>_____ |

## Lab X: End of lab

| Step | Action |
|---|---|
| | If there is spare time, you are welcome to re-do any of these exercises, or attempt to come up with your own activities in the environment. This is entirely up to you.<br><br>Before you leave, please clean up the lab environment for the next student assigned your account. |
| 1. | In From the main vCloud Director screen → Home page, click on the stop button on your vApp<br><br> 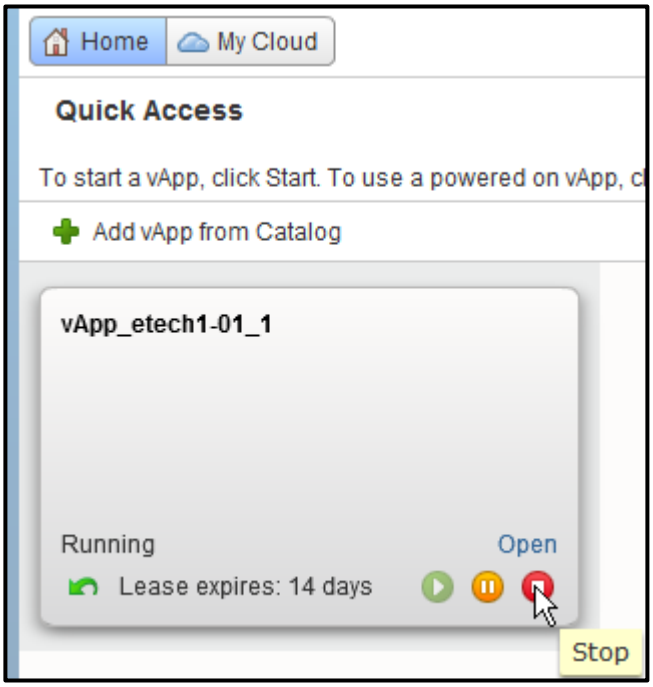 <br> |

| Step | Action |
|---|---|
| 2. | Wait for it to go from stopping to stopped, then right-click the vApp and click Delete 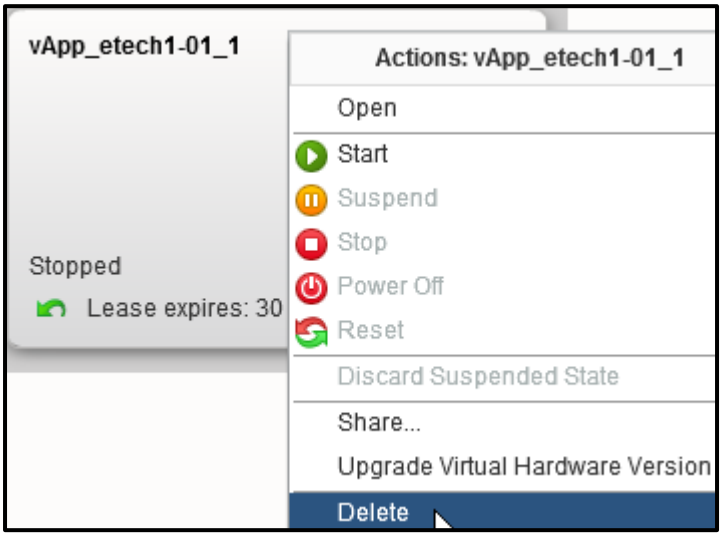  |
| 3. | Click Yes to proceed 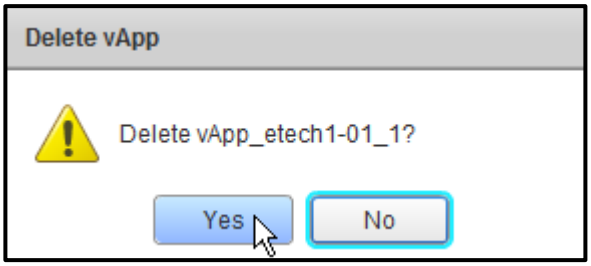  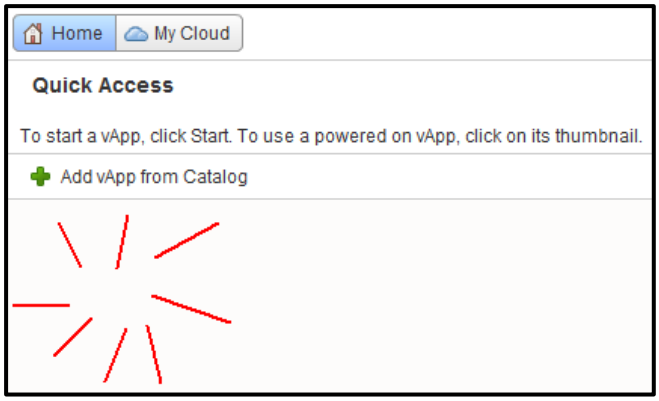  |

*End of Lab*