Dell EMC VxFlex Ready Node

Version 2.x

AMS Troubleshooting Guide

P/N 302-004-650

REV 02

DELL EMC

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

Published June 2018

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

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

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

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

CONTENTS

Preface **5**

Chapter 1 **Troubleshooting VxFlex Ready Node** **7**
AMS deployment fails during management IP configuration....8
AMS cannot access iDRAC/denial of service....8
AMS Deployment Wizard aborted.... 9
Application server does not see a VxFlex OS volume....9
Deployment fails at Add Device phase....10
Disk appears in Unconfigured Bad state....11
Entering ESX maintenance mode—best practice.... 11
Finding the physical location of a NIC port on an ESXi server.... 12
Issuing SCLI commands without the MDM password....13
Recovering an AMS server....14
Resolving failed ESX ISO image installation....16
Restarting the AMS service on Linux.... 16
VxFlex Ready Node CLI or GUI cannot connect to the AMS server....16
SDS devices are in “failed” state after an SVM crash.... 17
Solving VxFlex OS performance issues....17
The BMC (iDRAC) does not respond to pings.... 17
Mismatch in I/O counters.... 18
Improving AMS memory usage or refresh time by modifying AMS sampling delay.... 18
Speeding up rebuild and rebalance processes.... 19
Remove node from Backend is not supported....19
During Set Management IPs stage BMC Gateway not configured.... 19
Node type is unknown after restarting node 20
PM module does not load during a DTK ISO installation....20
VxFlex OS AMS service does not start after restart....21
NVDIMM error prevents boot....22
Toshiba PM4 SSD disks may fail when going into power-saving sleep....22
How to change iDRAC IP address for AMS managed node....23
Firmware update via AMS may fail during upload.... 25

Chapter 2 **VxFlex Ready Node Upgrade Troubleshooting** **27**
After retrying upgrade "Failed to query node" message is displayed.... 28
Upgrade freezes during SDS exit maintenance mode process....28
After upgrading, a "Could not switch to cluster mode" error message is displayed.... 28
After aborting an upgrade, the upgrade cannot be resumed....29
MDM package is not updated during upgrade.... 30
Upgrade requires ESX to enter maintenance mode.... 30
Upgrade fails when NIC in cluster node is down....31
Abort AMS upgrade causes lost connection to SDS....31

Preface

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

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

Note

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

Previous versions of Dell EMC VxFlex OS were marketed under the name Dell EMC ScaleIO.

Similarly, previous versions of Dell EMC VxFlex Ready Node were marketed under the name Dell EMC ScaleIO Ready Node.

References to the old names in the product, documentation, or software, etc. will change over time.

Note

Software and technical aspects apply equally, regardless of the branding of the product.

Related documentation

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

The following Dell EMC publication sets provide information about your VxFlex OS or VxFlex Ready Node product:

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

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

Typographical conventions

Dell EMC uses the following type style conventions in this document:

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)

Italic	Used for full titles of publications referenced in text
`Monospace`	Used for: • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
Monospace italic	Used for variables
`Monospace bold`	Used for user input
[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

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

Product information

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

Technical support

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

Your comments

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

CHAPTER 1

Troubleshooting VxFlex Ready Node

The following topics include troubleshooting events and suggested solutions related to the VxFlex Ready Node system.

- AMS deployment fails during management IP configuration 8
- AMS cannot access iDRAC/denial of service 8
- AMS Deployment Wizard aborted 9
- Application server does not see a VxFlex OS volume 9
- Deployment fails at Add Device phase 10
- Disk appears in Unconfigured Bad state 11
- Entering ESX maintenance mode—best practice 11
- Finding the physical location of a NIC port on an ESXi server 12
- Issuing SCLI commands without the MDM password 13
- Recovering an AMS server 14
- Resolving failed ESX ISO image installation 16
- Restarting the AMS service on Linux 16
- VxFlex Ready Node CLI or GUI cannot connect to the AMS server 16
- SDS devices are in “failed” state after an SVM crash 17
- Solving VxFlex OS performance issues 17
- The BMC (iDRAC) does not respond to pings 17
- Mismatch in I/O counters 18
- Improving AMS memory usage or refresh time by modifying AMS sampling delay 18
- Speeding up rebuild and rebalance processes 19
- Remove node from Backend is not supported 19
- During Set Management IPs stage BMC Gateway not configured 19
- Node type is unknown after restarting node 20
- PM module does not load during a DTK ISO installation 20
- VxFlex OS AMS service does not start after restart 21
- NVDIMM error prevents boot 22
- Toshiba PM4 SSD disks may fail when going into power-saving sleep 22
- How to change iDRAC IP address for AMS managed node 23
- Firmware update via AMS may fail during upload 25

AMS deployment fails during management IP configuration

When the AMS deployment process fails during management IP configuration, and the system displays an error message as `Cannot access the IPs (xxxx) on the same management network due to: The node is unreachable`, the likely cause of this issue is network related.

The following table describes the likely causes of deployment issues and their recommended solutions:

Table 1 Causes of AMS deployment issues

	Root Cause	Action
1.	Configured incorrect VLAN ID, during initial wizard	Refer to the Network section in VxFlex OS User Guide
2.	Configured different layer-2 networks for the BMC (iDRAC) IP setup and the AMS network.	Refer to the Network section in VxFlex OS User Guide
3.	AMS firewall ports in blocked state	Refer to the Security Guide

Procedure

1. During AMS deployment, if an error message appears, identify the likely cause of failure.

   The likely cause of failure is network related, as described in the table above.
2. Resolve the network-related issue and click **Retry** to continue with the process of AMS deployment.
3. Do the following:
   - If the AMS deployment process resumes, allow the procedure to continue, until it is completed.
   - If the AMS deployment issue is not resolved, go to the next step.
4. Remove the node and go back to the factory reset configuration.
5. Perform the procedure of adding a node.
6. Complete the procedure of management IP configuration and AMS deployment.

AMS cannot access iDRAC/denial of service

AMS access to the iDRAC over the web is locked, and the iDRAC cannot be opened with a web browser from the AMS server.

Before you begin

Ensure that you have the following:

- Root access to all nodes in the system.
- Product ISO.

If the iDRAC password was manually changed before properly updating the password in the AMS, the AMS sampling can shut down iDRAC access from the AMS IP address.
The solution depends on the scenario in which the issue occurs:

- During proof of concept - Reinstall the AMS and reimage the node.
  See the *VxFlex Ready Node AMS Deployment Guide*.
- Password change on existing nodes - follow the procedure below.

Procedure

1. Stop the AMS service.
2. Back up all the AMS configuration files:
   - Linux:

     ```
     cd /opt/emc/scaleio/ams; tar cf - cfg webapps/ROOT/WEB-INF/
     classes ) | gzip -c > "/tmp/ams_upgrade_$(date +'%Y%m%d%H%M
     %S').tar.gz"
     ```

     The `/tmp/ams_upgrade_<current_date_time>.tar.gz` backup file is created.
   - Windows:

     ```
     mkdir C:\Windows\Temp\AMS_Backup\cfg & mkdir C:\Windows\Temp
     \AMS_Backup\classes & xcopy /e/c/k/y "C:\Program Files\EMC
     \scaleio\AMS\cfg" "C:\Windows\Temp\AMS_Backup\cfg" &
     xcopy /e/c/k/y "C:\Program Files\EMC\scaleio\AMS\webapps
     \ROOT\WEB-INF\classes" "C:\Windows\Temp\AMS_Backup\classes"
     ```

     The `C:\Windows\Temp\AMS_Backup\` backup directory is created.
3. Remove the AMS package from the server.
4. Change all the iDRAC passwords.
5. Install the AMS again.
6. Run the AMS recovery flow, as described in the "Recovering an AMS server" section.

AMS Deployment Wizard aborted

Aborting the AMS Deployment Wizard does not allow you to retry the deployment.

When you select to abort at any stage during the AMS Deployment Wizard, there may be no option to retry the deployment. To redeploy, perform the following:

Procedure

1. Close the AMS Deployment Wizard.
2. Remove the nodes from the system.

   For more information, see "Removing a node from the system" in the user documentation.
3. Perform factory reset to the nodes. as described in Performing a factory reset in the *VxFlex Ready Node AMS User Guide*
4. Run the AMS Deployment Wizard again.

Application server does not see a VxFlex OS volume

Perform the following steps:

Procedure

1. Check if the VxFlex OS system is operational:

```
scli --mdm_ip <mdm IP> --query_all
```

2. Check if the volume is mapped to any of the SDC servers:

```
--query_all_volumes
```

3. Determine if the SDC is installed on the server:
   - Linux: Run `rpm -qa | grep sdc`
   - ESX: Run `esxcli software vib list|grep sdc`
4. Determine if the SDC is connected to an MDM:

```
scli --mdm_ip <mdm IP> --query_all_sdc
```

5. Ensure that the MDM management IP address is up and running.
6. On an application server, rescan for new volumes by performing one of the following:

Operating system	Description
RHEL	Run `fdisk -l | grep scini` to see the volume list.
ESXi	Perform one of the following: • Rescan for new devices, using the vSphere Rescan adapter. • Run `esxcli storage core device list|grep eui`

Deployment fails at Add Device phase

If a storage device was previously used for a different purpose, and was partitioned, deployment may fail. To rectify this problem, follow these steps:

Procedure

1. Leave the **Deployment** window open.
2. In command line, to find the name of the partitioned storage device, type the command `cat /proc/partitions`

   Check the output for partitioned devices, and make a note of their names. For example, in the following output, the device called `sdf` has three partitions.

```
scaleiovm:/opt/emc/scaleio/mdm/bin # cat /proc/partitions
major minor #blocks name
```

```
8 48 878542848 sdd
8 64 878542848 sde
8 80 878542848 sdf
8 81 4064 sdf1
8 82 4193280 sdf2
8 83 870806511 sdf3
```

3. Type the command `hdparm -z /dev/<name_of_partitioned_device>`

   where <name_of_partitioned_device> is the device name identified in the previous step.

   The partitions have been removed. You can issue the `cat /proc/partitions` command again to verify this.

4. In the **Deployment** window, click **Retry** and continue the deployment process.

Disk appears in Unconfigured Bad state

Removing, and then replacing a disk in a chassis will change its state to `Unconfigured Bad`.

To clear this state, perform the following commands from the SVM:

Procedure

1. Identify the disk information by running this command:

   ```
   /opt/MegaRAID/perccli/perccli64 /c0/eall/sall show
   ```

   From the output, identify the enclosure ID (EID) and slot number (Slt) of the failed disk.

2. Run `/opt/MegaRAID/perccli/perccli64 /c0/e<EID>/s<Slt> set good`, where *<EID>* and *<Slt>*use the output from the previous step.
3. Run `/opt/MegaRAID/perccli/perccli64 /c0/ /fall import`

   The disk state should now return to normal.

Entering ESX maintenance mode—best practice

If you do not shut down a Storage VM (SVM) before entering vSphere maintenance mode, maintenance mode will hang for a long time while the shared Datastores are moved to other VMs, and the VM will not shut down automatically. To prevent this problem, perform the following steps:

Procedure

1. Before entering maintenance mode, use the VxFlex OS GUI to check that there is no rebuild or rebalance operation in progress. If indeed there is no such operation, shut down the SVM. This will initiate a rebuild flow.
2. After entering maintenance mode on one node, ensure that this node's rebuild completes before entering maintenance mode on another node.
3. When you exit maintenance mode, power on the SVM. This will initiate a rebalance operation. Ensure that it is completed successfully before performing any other operations.

Finding the physical location of a NIC port on an ESXi server

This topic explains how to find the physical location of a NIC port on a node, using its ESX logical PCI slot location.

To find the NIC port, perform the following steps:

Procedure

1. Note the logical slot address shown for the port in vSphere. The logical PCI slot address is represented by `vmnicX XXXX:XXX:XX:X`. For example: `vmnic2 0000:082:00.0`. You can also generate a list of all the logical PCI slot addresses, using the command:

   ```
   esxcli network nic list
   ```

   Output similar to the following will be displayed:

   ```
   Name PCI Device Driver Link Speed Duplex MAC Address MTU
   Description
   ------ ------------- ------ ---- ----- ------
   ----------------- ----
   ---------------------------------------------------------------
   -
   vmnic0 0000:005:00.0 ixgbe Down 0 Half 2c:60:0c:45:3f:bc 1500
   Intel Corporation 82599 10 Gigabit Dual Port Network
   Connection
   vmnic1 0000:005:00.1 ixgbe Down 0 Half 2c:60:0c:45:3f:bd 1500
   Intel Corporation 82599 10 Gigabit Dual Port Network
   Connection
   vmnic2 0000:082:00.0 ixgbe Up 10000 Full 90:e2:ba:82:be:34
   1500 Intel Corporation 82599 10 Gigabit Dual Port Network
   Connection
   vmnic3 0000:082:00.1 ixgbe Up 10000 Full 90:e2:ba:82:be:35
   1500 Intel Corporation 82599 10 Gigabit Dual Port Network
   Connection
   vmnic4 0000:001:00.0 igb Up 1000 Full 2c:60:0c:73:01:b0 1500
   Intel Corporation I350 Gigabit Network Connection
   vmnic5 0000:001:00.1 igb Down 0 Half 2c:60:0c:73:01:b1 1500
   Intel Corporation I350 Gigabit Network Connection
   ```

2. Once you know the logical PCI address, type the following command to display more information about the logical address:

   ```
   esxcli hardware pci list
   ```

   Output similar to the following is displayed. The first line in the output below (`000:082:00.0`) represents the logical PCI slot location. Find the information for the required logical PCI slot. The lines `Physical Slot` and `Slot Description` represent the physical slot location.

   ```
   000:082:00.0
   Address: 000:082:00.0
   Segment: 0x0000
   Bus: 0x82
   Slot: 0x00
   Function: 0x00
   ```

```
VMkernel Name: vmnic2
Vendor Name: Intel Corporation
Device Name: 82599 10 Gigabit Dual Port Network Connection
Configured Owner: Unknown
Current Owner: VMkernel
Vendor ID: 0x8086
Device ID: 0x10fb
SubVendor ID: 0x8086
SubDevice ID: 0x0003
Device Class: 0x0200
Device Class Name: Ethernet controller
Programming Interface: 0x00
Revision ID: 0x01
Interrupt Line: 0x0b
IRQ: 11
Interrupt Vector: 0x37
PCI Pin: 0x00
Spawned Bus: 0x00
Flags: 0x0201
Module ID: 4124
Module Name: ixgbe
Chassis: 0
Physical Slot: 8
Slot Description: Chassis slot 8.00
Passthru Capable: true
Parent Device: PCI 0:128:3:0
Dependent Device: PCI 0:130:0:0
Reset Method: Function reset
FPT Sharable: true
```

Issuing SCLI commands without the MDM password

This topic describes how to find the Master MDM and how to reset its credentials so that you can run VxFlex OS scli commands.

Performing the procedure requires write privileges on the Master MDM.

Use this procedure when you need to run VxFlex OS SCLI commands, but you do not have the MDM credentials.

Procedure

1. Find the two nodes that are acting as MDM servers, and determine which is the Master MDM.

   a. Using a text editor, open the `ams.cfg` file on the AMS server.

   - Windows: `C:\Windows\System32\config\systemprofile\AppData\Roaming\emc\scaleio\ams\cfg\nodes.cfg`
   - Linux: `/opt/emc/scaleio/ams/cfg/`

   b. Determine the IP addresses for the MDM servers for which `scaleioRole="CLUSTER_ROLE_MDM"`.

   Example:

```
"Node-id-1-4172547358935" : {
"uuid" : "Node-id-1-4172547358935",
"isStale" : null,
"nodeProfile" : null,
"osType" : null,
"scaleioRole" : "CLUSTER_ROLE_MDM",
"hostManagementIp" : "10.0.0.2",
```

```
"svmManagementIp" : "10.0.0.12",
"bmcIp" : "10.0.1.2",
"hostPrimaryDataIp" : "20.0.0.2",
"hostSecondaryDataIp" : "20.0.1.2",
"svmPrimaryDataIp" : "20.0.0.7",
"svmSecondaryDataIp" : "20.0.1.7",
```

2. SSH to the svmManagementIp address of each MDM server, and type `scli --query_cluster`. The Master MDM will return a genuine response, while a Slave MDM will return an error.
3. Reset the password for the admin user:
   a. On the Master MDM, create a text file called `MDM_SERVICE_MODE`, in the location corresponding to your operating system:
      - Windows: `C:\Program Files\emc\scaleio\MDM\logs`
      - Linux: `/opt/emc/scaleio/mdm/logs/`
   b. In the body of the file, type the text `Reset Admin`, and save the file.
   c. From the MDM, type `scli --reset_admin`.

      The admin user's password is reset to admin.

Results

Now that you have the Master MDM IP, user name (admin), and password (admin), you can issue scli commands.

Recovering an AMS server

If an AMS server goes down, or if you want to replace an AMS server, you can run the system recovery process to recover the configuration details for all of its nodes, including nodes that are disconnected. Additionally, if nodes are in the process of being deployed when the AMS goes down, this information is recovered and the nodes can continue with the deployment once the AMS server has been recovered. vCenter and ESRS information are also recovered.

Before you begin

Ensure that:

- You have downloaded the AMS and VxFlex OS GUI packages.
- You install the same AMS build on the new AMS server as was installed on the one you are replacing.
- You save the `phoenix.cfg`, `phoenix_bak.cfg` and `phoenix.properties` files from the old AMS server. The file locations are:
  - `phoenix.cfg` and `phoenix_bak.cfg`:
    Windows: `C:\Program Files\EMC\scaleio\AMS\cfg`

    Linux: `/opt/emc/scaleio/AMS/cfg`
  - `phoenix.properties`:
    Windows: `C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\WEB-INF\classes`

    Linux: `/opt/emc/scaleio/ams/webapps/ROOT/WEB-INF/classes`
- Optionally, that you have prepared a file containing node credentials, using the template located in:

- *Windows*: `C:\Program Files\EMC\scaleio\AMS\webapps\ROOT\resources\RecoveryUserDataTemplate.csv`
- *Linux*: `/opt/emc/scaleio/ams/webapps/ROOT/WEB-INF/resources/RecoveryUserDataTemplate.csv`

Procedure

1. Install the AMS and VxFlex OS GUI packages on a new server.
2. Open the VxFlex OS GUI and log in to the new AMS server.
3. Set the initial AMS password.
4. Perform the initial AMS setup procedure, as described in the *VxFlex Ready Node Deployment Guide*, to enter the initial IP address information and other credentials.

   Note

   In the **Security** screen, you must use the same default passwords that were used in the AMS server being replaced.

5. At the end of the initial setup, when the **Add Nodes** option is displayed, close the initial AMS setup wizard.
6. From a command prompt, change the working directory to:
   - Linux-based AMS: `/opt/emc/scaleio/ams/webapps/ROOT/resources/scaleio_repository/`
   - Windows-based AMS: `\Program Files\EMC\scaleio\AMS\webapps\ROOT\resources\scaleio_repository`
7. At the command prompt, run the amscli command to recover the system (square brackets indicate the parameter is optional):

   ```
   java -jar EMC-ScaleIO-ams-cli-2.6-<build>.jar --username
   <USERNAME> --password <PASSWORD> --recover_system [--
   ssh_keys_file <SSH_FILE>] [--custom_node_credentials_file
   <CREDENTIALS_FILE>]
   ```

   where:

   - *<USERNAME>* is the current AMS username.
   - *<PASSWORD>* is the current AMS password.
   - --ssh_keys_file (optional) indicates that the file at *<SSH_FILE>* contains the SSH keys to the nodes being recovered and that these should be applied to the nodes.

     Note

     The SSH keys file can be downloaded using the AMS. In the VxFlex OS GUI, open the **System Settings** and select **Connection**. Click **Backup Keys**.

   - --custom_node_credentials_file (optional) indicates that the file at *<CREDENTIALS_FILE>* contains the node credentials.

8. Log in to the new AMS and verify that the node configuration has been recovered.
9. Copy the `phoenix*.*` files back to the same locations on the AMS server as mentioned in the Prerequisites section at the beginning of this topic.

10. Restart the AMS service.

Resolving failed ESX ISO image installation

When you face an unexpected failure of ESX ISO image installation, with error message as `Cannot find kickstart file on cd-rom with path - /KS/KS,CFG`), the recommended workaround is:

- Map the ISO image again and reinstall the ESX.

Restarting the AMS service on Linux

This topic explains how to restart the AMS service on Linux, when the usual procedures do not succeed.

To restart the AMS service, perform the following steps:

Procedure

1. Type the command `service scaleio-ams stop`.
2. Stop all processes manually, using either the `killall` command, or by PIDs.

   For example:

   - ```
     ps -ef | grep java
     ```
   - ```
     kill -9 PID
     ```

3. Start the AMS service again, by typing the command:

   ```
   service scaleio-ams start
   ```

VxFlex Ready Node CLI or GUI cannot connect to the AMS server

Perform the following steps:

Procedure

1. Ensure that you are connecting to the IP address of the AMS server.
2. Ping the AMS IP address to ensure you have connectivity.
3. Check if the AMS service is running:

   a. Windows: In the **Services** window, check the **EMC ScaleIO AMS** service.

   b. Linux : Run the following commands:

   ```
   chkconfig --list | grep -i scaleio
   ```

   ```
   ps -ef | grep -i tomcat
   ```

Note

If the AMS was installed as a proxy server, its name will be slightly different.

4. Ensure that the management IP address is up and running.

SDS devices are in “failed” state after an SVM crash

If the SVM crashes, when it comes back up, a rebuild is initiated, and all SDS devices may be in "failed" state.

To fix this problem, restart the ESX server.

Solving VxFlex OS performance issues

VxFlex OS is designed to generate the best performance possible from any given system configuration, by using all possible nodes and distributing the data evenly among them.

When the system performance does not meet your expectations, verify if:

- The relevant volume is allocated to a high-performance Storage Pool.
  For example, is it allocated to a pool consisting of SSDs only? If not, using such a Storage Pool will generate better performance.
- The relevant volume resides in a Storage Pool that consists of different storage drives, with different performances?
  A low performance drive in the pool will slow down all the members (waiting for it to respond). If possible, avoid mixing different types of drives.
- The network in use provides maximal network bandwidth to all the ports in use by VxFlex OS.

For a full performance review, use the following resources:

- For all VxFlex OS-related products, see *VxFlex OS Performance Fine-Tuning Technical Notes*.
- For VxFlex OS software only, see *System Analysis Best Practice Guide*.

After following these suggestions, you may contact EMC Support for professional analysis and assistance.

The BMC (iDRAC) does not respond to pings

In unlikely event that the BMC (iDRAC) was configured correctly by the deployment process, but does not respond to pings, perform the following workaround:

Procedure

1. Open an SSH session with the ESX’s IP address that was set during the **Set Management IP Addresses** stage, using the credentials set during the initial configuration with the wizard.
2. Type the following command:

```
/vmfs/volumes/scaleio-datastore/IPMITOOL mc reset cold
```

After 20–30 seconds, the BMC (iDRAC) module starts responding.

Mismatch in I/O counters

Resolving mismatch in the I/O counters of the VxFlex OS Dashboard and customer application.

The I/O counter in the VxFlex OS Dashboard might display a read/write I/O value which is larger than the one displayed in the customer application, although the total bandwidth in both the cases is the same.

The mismatch between the I/O counters occurs when the customer operating system, which triggers the I/O requests to the SDS, splits them into two or more smaller-sized I/O requests; the total bandwidth being the same as the bandwidth of the original I/O request. Thus the I/Os are split at the level of the customer operating system (which causes VxFlex OS to see a larger number of I/Os), and are integrated as well at the same level (the level that splits them).

Currently, VxFlex OS does not support any workaround for such scenarios.

Improving AMS memory usage or refresh time by modifying AMS sampling delay

AMS memory usage and VxFlex OS GUI refresh are influenced by the quantity of nodes managed by the AMS. Modify the sampling delay to suit your system setup.

If the AMS manages a large number of nodes, increasing the sampling delay between queries will positively impact on AMS memory usage. If the AMS manages a small number of nodes, reducing the sampling delay will provide more frequent refreshes of hardware-related items shown in the VxFlex OS GUI.

Procedure

1. Open the `ams.properties` file in a text editor. The file is located in the following location on the AMS server:

Operating System	Path
Linux	`EMC\ScaleIO\AMS\webapps\ROOT\WEB-INF\classes\ams.properties`
Windows	`C:\Program Files\EMC\ScaleIO\AMS\webapps\ROOT\WEB-INF\classes\ams.properties`

2. Edit the value for the property *samplingDelay*, and then save the file.

   This value sets the time (in seconds) between the end of one sampling cycle, and the beginning of the next one (per node).

3. Restart the AMS service.

Results

The sampling delay has been modified.

Speeding up rebuild and rebalance processes

You can set concurrent activity limits to speed up rebuild and rebalance processes.

You can speed up rebuild and rebalance processes. This can be very useful If you are planning maintenance, have a disaster recovery (DR) situation, or must have faster rebuild/rebalance times.

You can use the SCLI or the VxFlex OS GUI:

- SCLI - Use the `scli --set_rebuild_rebalance_parallelism` command to increase the amount of concurrent activities.
- VxFlex OS GUI - From the **Backend** view, sort by Storage Pool. Right-click the Storage Pool and select **Set I/O Priority**.

Check the service-level agreement (SLA) of your environment/application to ensure that the selected setting does not adversely affect your applications/clients. Test your environment to determine what the optimal "limit" is.

The following example sets concurrent activity to 10:

```
scli --set_rebuild_rebalance_parallelism
--protection_domain_name default --storage_pool_name default
--limit 10 --mdm_ip 10.13.168.138
```

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

Remove node from Backend is not supported

Removing a node from the`Backend` view is not supported.

Scenario:

1. Rack 1 is decommissioned and needs to be replaced. It holds the physical AMS.
2. A new AMS is installed.
3. Recovery for the new nodes and a new AMS server is initiated.
4. Missing nodes are not identified and one of the nodes is the MDM.
5. Faulty nodes are removed
6. The MDM role is migrated to a new node.

Expected:

Remove node is not supported from the`Backend` view. Migrating MDM to a new node is successful.

Actual:

Removing a node from the`Backend` view is not supported. Migration fails because the MDM is not in the list since it was not detected during recovery.

During Set Management IPs stage BMC Gateway not configured

During the Set Management IPs stage, the BMC Gateway was not configured.

Problem:

1. During the `Set Management IPs` stage, all nodes failed with error:

```
Command failed: Node ESXSystemNode[nodeName=<name>,nodeIPs=<ips>,
credential=<null>] cannot access other IPs (<bmc_ip>) on the
same data network due to 2
```

2. Tried to ping the BMC from ESXi. There was no response.
3. Ran the IPMI command on the ESXi:`#/vmfs/volumes/scaleio-datastore/IPMITOOL lan print`

```
Set in Progress : Set Complete Auth Type Support : NONE MD2 MD5
PASSWORD OEM Auth Type Enable : Callback : NONE MD2 MD5
PASSWORD OEM : User : MD2 MD5 PASSWORD OEM : Operator : MD2 MD5
PASSWORD OEM : Admin : MD2 MD5 PASSWORD OEM : OEM : IP Address
Source : Static Address IP Address : 10.103.109.71 Subnet
Mask : 255.255.252.0 MAC Address : 2c:60:0c:06:b4:d0 SNMP
Community String : Quanta IP Header : TTL=0x40 Flags=0x40
Precedence=0x00 TOS=0x10 BMC ARP Control : ARP Responses
Enabled, Gratuitous ARP Disabled Gratituous ARP Intrvl : 0.0
seconds Default Gateway IP : 0.0.0.0 Default Gateway MAC :
00:00:00:00:00:00 Backup Gateway IP : 0.0.0.0 Backup Gateway
MAC : 00:00:00:00:00:00 802.1q VLAN ID : Disabled 802.1q VLAN
Priority : 0 RMCP+ Cipher Suites : 0,1,2,3,6,7,8,11,12 Cipher
Suite Priv Max : caaaXXaaaXXaaXX : X=Cipher Suite Unused :
c=CALLBACK : u=USER : o=OPERATOR : a=ADMIN : O=OEM
```

The results show that the Gateway was not configured

Solution:

- Run the following command via the ESXi:`/vmfs/volumes/scaleio-datastore/IPMITOOL lan 1 set defgw ipaddr <Gateway IP>`
  The command sets the Gateway and continues with the deployment.

Node type is unknown after restarting node

After restarting a node, the message `The node type is unknown` is displayed in the **Hardware** view.

Problem:

1. Deployed three bare-metal node systems (RH7.3/4).
2. Restarted one of the nodes.
3. In the **Hardware** view, the following message is displayed:`The node type is unknown`

Solution:

Run the AMS CLI query:`--query_all_node`

After several minutes, the node appears in the **Hardware** view.

PM module does not load during a DTK ISO installation

Problem:

During the installation of the DTK ISO, the PM module does not load.

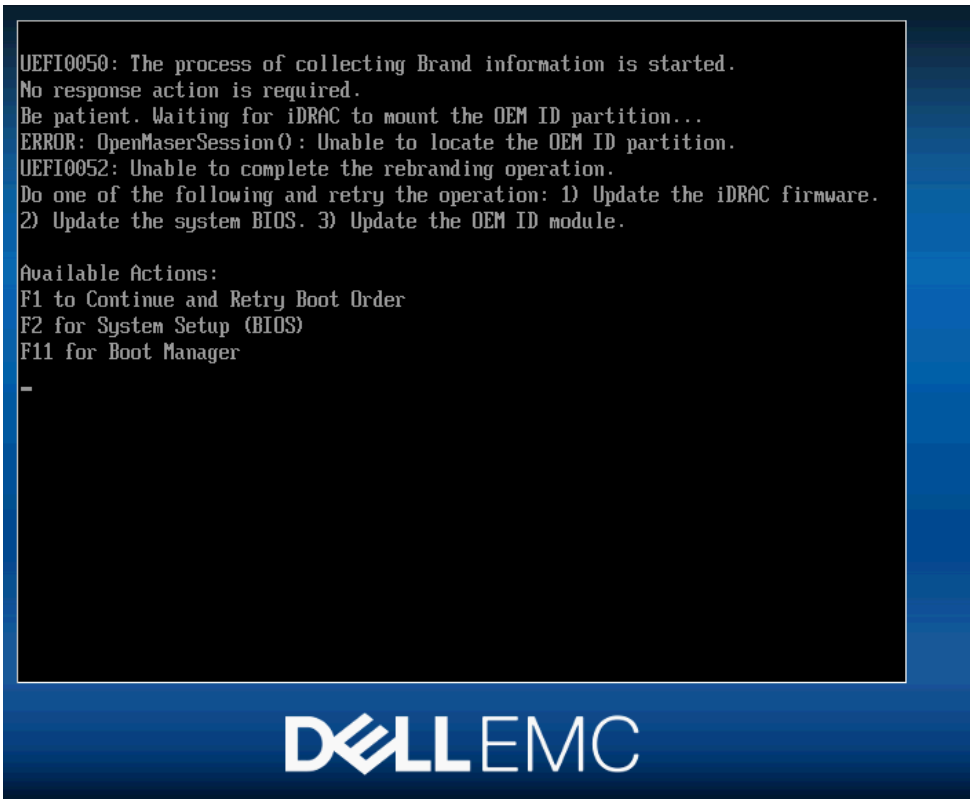

Solution:

Note

This is relevant after an upgrade.

1. From the iDRAC window, select **Maintenance** > **Diagnostics**.

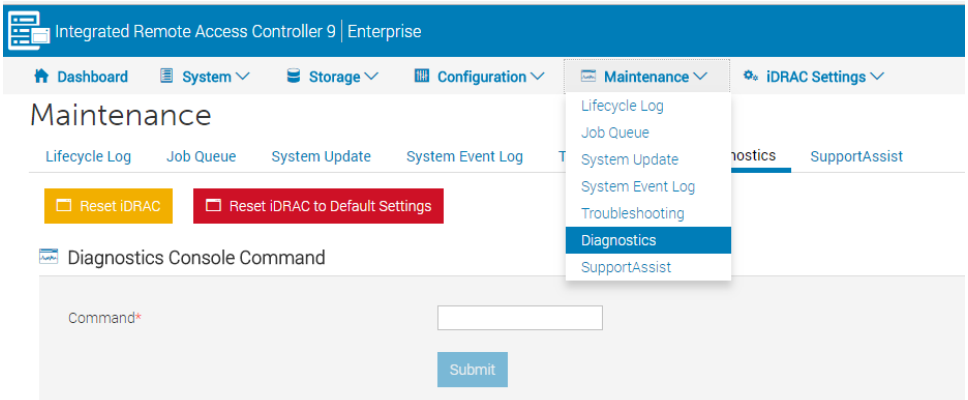

2. Click **Reset iDRAC** and follow the online instructions.
3. Restart the node.

VxFlex OS AMS service does not start after restart

Note

This is relevant to Windows AMS

Problem:

In some extreme cases VxFlex OS AMS service does not start after restart of the AMS server or manually from Windows Services.

Solution:

Rerun the MSI of AMS to repair the AMS component and enable the service to start properly.

NVDIMM error prevents boot

Problem:

The following error prevents the server to boot:

```
UEFI0302:
Save Error on NVDIMM-N located at B7. This NVDIMM-N module will be
set to read only
mode.
```

Solution:

First option:

1. Go to System Setup (BIOS).
2. Open **Memory Settings** > **Persistent Memory**.
3. Scroll down to **NVDIMM** and click **Enabled** to set to factory defaults.
4. Save settings and restart server.

Second option:

1. Remove input power to the system.
2. Reset the NVDIMM-N module.
3. Restart the server.

If issue still persists, replace the faulty module identified in the message.

Toshiba PM4 SSD disks may fail when going into power-saving sleep

Problem:

Toshiba PM4 SSD disks may erratically fail when going into power-saving sleep. These disks have a lockup mechanism that causes VxFlex OS to timeout on non-responsive disk.

Note

If using ESX/Windows, iDRAC is the preferred method. Instructions can be found on the Dell site at: https://www.dell.com/support/home/us/en/19/drivers/driversdetails?driverId=0RTJ1

Solution:

To upgrade the devices firmware, you can obtain the firmware (.BIN file) from the device manufacturer or from the Dell support site. The software file required for upgrade is in the format "SAS-Drive_Firmware_XXXX_.BIN"

Perform the following steps on every SDS node that contains a device that requires the upgrade:

1. Copy the package to the `/tmp` folder.
2. Use SSH to connect to the SDS node as root.
3. Make the BIN file executable: run `chmod +x /tmp/SAS-Drive_Firmware_XXXX_.bin`
4. Using the VxFlex OS CLI or the VxFlex OS GUI, enter the SDS node into maintenance mode (either run "`scli --enter_maintenance_mode --sds_name <sds_name>` command or from the VxFlex OS GUI select SDS node and right-click on **Enter maintenance mode**. Wait until all IO on SDS devices is stopped.
5. Validate that the SDS is in maintenance mode: marked in maintenance mode in VxFlex OS GUI or run `scli --query_all_sds` and look for the selected SDS **State**.
6. Stop the SDS process. On the SDS, run this command, as root: `/opt/emc/scaleio/sds/bin/delete_service.sh`.
7. Execute `FW BIN file: /tmp/SAS-Drive_Firmware_XXXX_.BIN`. Optionally, you may execute it with `-q` option to prevent further need to approve SW agreement and actions (quite mode).
8. Follow the installation instructions and wait for the upgrade to complete. DO NOT interrupt the process, restart the machine and install any other program during the process.
9. After completion, start the SDS process: `/opt/emc/scaleio/sds/bin/create_service.sh`
10. Exit maintenance mode: either using VxFlex OS GUI (right-click on **SDS** > **Exit maintenance mode**") or scli (run `scli --exit_maintenance_mode --sds_name <sds_name>`
11. Wait for **rebuild/rebalance** action to complete.
12. Repeat all of these steps on the next SDS in the Storage Pool, and then on every other SDS in the Storage Pool that has devices that must be upgraded.
13. To verify the firmware was applied to the disks, run this command (assuming PERCCLI is installed): `/opt/MegaRAID/perccli/perccli64 /c0/eall/sall show all |grep "Firmware Rev" -B7 |grep -v "Man|NAND|Mod|WWN`
14. You can also verify from iDRAC: **Storage** -> **Physical Disks** and then click the + symbol next to each disk and look at **Revision**.

How to change iDRAC IP address for AMS managed node

This article describes how to change an AMS-managed VxFlex Ready Node Baseboard Management Controller's (BMC/iDRAC) IP address for a node that is already part of a VxFlex OS system.

Solution

There are two options for resolving the issue.

Recommended method:

1. From the VxFlex OS GUI, in the **Hardware** view, right-click the node and select **Remove VxFlex OS**.
   In the **Remove VxFlex OS** dialog, click **OK**.

2. From the VxFlex OS GUI, in the **Hardware** view, right-click the node and select **Remove Node & Unassign IP Addresses**.
   In the **Remove Node & Unassign IP Addresses** dialog, click **OK**.
3. In the **Management Network** panel, from **System Settings** > **Management and Data IPs** tab, change the IP address range for **iDRAC**.
4. Perform a factory reset of the node as described in *VxFlex Ready Node AMS User Guide*.
5. From the **Hardware Deployment Wizard**, deploy the new node.
6. During the deployment, assign the preferred iDRAC IP address to the node.

Alternative method:

1. Change the iDRAC IP address as described in *VxFlex Ready Node System Deployment Guide*.
   The iDRAC may appear as **Disconnected** in the AMS GUI.
2. In the **Management Network** panel, from the **System Settings** > **Management and Data IPs** tab, change the IP address range for **iDRAC**.
3. On the AMS host, stop the Dell EMC VxFlex OS AMS service.
4. Perform a backup of the AMS configuration files from the Linux server:

```
( cd /opt/emc/scaleio/ams; tar cf - webapps/ROOT/WEB-INF/
classes/ conf/ cfg/ |gzip -c > ~/"ams_conf_$(date +'%Y%m%d%H%M
%S').tar.gz" )
```

5. Open the `nodes.cfg` file for editing.
   The file is located at:
   - Linux - `%PROGRAMFILES%\EMC\ScaleIO\AMS\cfg\`
   - Windows - `/opt/emc/scaleio/ams/cfg`
6. Replace the iDRAC IP address with the new IP address in the `bmcIp` field for the relevant node.
7. Change the value of the `supportChecksum` field to `false`.
8. Save the `nodes.cfg` file.
9. From the `cfg` directory on the relevant operating system (as noted step 5), delete the following files:
   - `/nodes.md5`
   - `/nodes_bak.cfg`
   - `/nodes_bak.md5`
   - `/model_cache/model.cfg`
   - `/model_cache/modelreferences.cfg`
10. On the AMS host, start the Dell EMC VxFlex OS AMS service.
    The AMS begins a polling cycle. Until the AMS completes its first polling cycle after starting, the VxFlex OS GUI **Hardware** tab may display the message `The node type is unknown`, or some other message that may not be accurate.

    The alert `Unable to connect to the monitoring agent (BMC)` might appear. The message disappears after 10-15 minutes.
11. In the AMS VxFlex OS GUI, remove any unnecessary IP address ranges.

If the issue still persists, replace the faulty module identified in the error message.

Firmware update via AMS may fail during upload

AMS uses the WSMAN library to communicate with the iDRAC entity of the node. During a firmware upgrade process that is used in Deployment (to align firmware to the approved list) and Upgrade, AMS uploads the firmware files to the node prior to running the upgrade command. In some cases WSMAN returns an invalid instance that causes the upload to fail.

Solution

There are two options for resolving the issue.

Option 1:

- Deployment: Perform iDRAC reset and after 2 minutes click **Retry** from the **Deployment** window to upload the firmware and continue the process
- Upgrade: perform iDRAC reset and after 2 minutes click **Retry** from the **Deployment** window to upload the firmware and continue te process.

Option 2 Upgrade: Alternative method if **Retry** in the Upgrade flow didn't work:

1. Run the latest ISO *VxFlex Ready Node PowerEdge 14G Hardware Update Bootable ISO* to update any firmware components to the correct firmware version. For instructions, see "DTK - Hardware Update Bootable ISO" in the *VxFlex OS v2.x Upgrade Guide* .
2. Login to the ESXi server or to the vCenter client.
3. Exit the ESXi host from maintenance mode as described in the *VxFlex OS User Guide*.
4. Login into the SVM Host and create a service of the SDS from the following path:

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

5. From the AMS GUI exit the SDS from maintenance mode as described in the *VxFlex OS User Guide*.
6. Click **Retry** from the **Upgrade** dialog, to continue the process.
   Repeat the steps on each of the nodes that failed.

Option 2 Deployment: Alternative method if **Retry** in the Deployment flow didn't work:

1. Run the latest ISO *VxFlex Ready Node PowerEdge 14G Hardware Update Bootable ISO* to update any firmware components to the correct firmware version. For instructions, see "DTK - Hardware Update Bootable ISO" in the *VxFlex OS v2.x Upgrade Guide* .
2. After the server boots up, click **Retry** from the **Deployment** window.
   Repeat the steps on each of the nodes that failed

CHAPTER 2

VxFlex Ready Node Upgrade Troubleshooting

The following topics include troubleshooting events and suggested solutions related to upgrading the VxFlex Ready Node system.

- After retrying upgrade "Failed to query node" message is displayed 28
- Upgrade freezes during SDS exit maintenance mode process 28
- After upgrading, a "Could not switch to cluster mode" error message is displayed 28
- After aborting an upgrade, the upgrade cannot be resumed 29
- MDM package is not updated during upgrade 30
- Upgrade requires ESX to enter maintenance mode 30
- Upgrade fails when NIC in cluster node is down 31
- Abort AMS upgrade causes lost connection to SDS 31

After retrying upgrade "Failed to query node" message is displayed

If upgrade via the AMS fails and you restart the upgrade using the VxFlex OS GUI **Retry** button, the following error message is returned: `Failed to query node`

Procedure

1. Stop the AMS service.
2. Delete the following file:

   `C:\Program Files\EMC\ScaleIO\AMS\cfg\asyncProcesses.cfg`
3. Start the AMS service.
4. Retry the upgrade.

Upgrade freezes during SDS exit maintenance mode process

NDU can be started when there is a faulty disk (failed capacity) in the Storage Pool. As a result, the upgrade process freezes when exiting the SDS from maintenance mode.

Procedure

1. Forcibly exit maintenance mode by running the following command on the AMS machine:

   ```
   mode:/opt/emc/scaleio/ams/webapps/ROOT/resources/scli --
   query_sds --sds_name <name>
   ```

2. In the VxFlex OS GUI **Upgrade** tab, click **Retry**.

Results

The upgrade proceeds.

After upgrading, a "Could not switch to cluster mode" error message is displayed

Upgrades using the VxFlex OS GUI might fail, and generate an error message similar to the following: `Command failed: Could not switch to cluster mode on X.X.X.X, X.X.X.X due to: Error`

Procedure

1. Restart the AMS service.
2. In the VxFlex OS GUI **Upgrade** tab, click **Close**.
3. Start the upgrade again.

After aborting an upgrade, the upgrade cannot be resumed

If an upgrade is manually aborted from the VxFlex OS GUI, the upgrade process cannot be resumed.

Procedure

1. Verify that the MDM process is running on all SVMs in the cluster:

```
ps -ef | grep -i mdm
```

2. If the MDM process is not running, start the MDM service:

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

3. Verify that the SDS process is running on all SVMs:

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

4. If the SDS process is not running, start the SDS service:

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

5. Navigate to the `scli` folder on the AMS:
   - Linux: `/opt/emc/scaleio/ams/webapps/ROOT/resources/`
   - Windows: `C:\Program Files\EMC\ScaleIO\AMS\webapps\ROOT\resources`
6. Check the cluster state:

```
./scli --query_cluster --mdm_port 6611
```

7. If the MDM cluster is in single-node mode, switch to 3-node mode:
   a. Log in to the MDM:

```
./scli --login --username <USERNAME> --password <PASSWORD>
```

   b. Allow commands during upgrade:

```
./scli --allow_commands_during_upgrade --mdm_port 6611
```

c. Switch to 3-node mode:

```
./scli --switch_cluster_mode --cluster_mode 3_node --
add_slave_mdm_ip <SLAVE_IP> --add_tb_ip <TB_IP> --mdm_port
6611
```

8. Restart the AMS service.
9. Use the VxFlex OS GUI to close the upgrade.
10. Use the VxFlex OS GUI to start the upgrade.

Results

The upgrade process begins.

MDM package is not updated during upgrade

After a successful upgrade, the MDM packages may not have been updated to the latest version.

If this happens, run the upgrade again. On the second attempt, the MDM artifacts will be upgraded.

Upgrade requires ESX to enter maintenance mode

During the upgrade process, your ESX may require to enter maintenance mode. If this is the case, and if it cannot do so, a message, similar to the following is displayed:

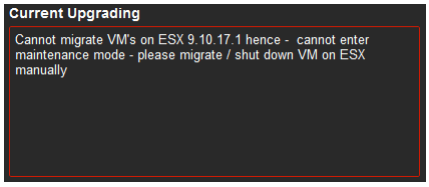

To continue, perform the following:

Procedure

1. In the vCenter client, find the ESX that is identified in the message, in our example, 9.10.17.1.
2. From the Virtual Machines tab, shut down the non- (Storage VM) on all the hosts that are being upgraded:

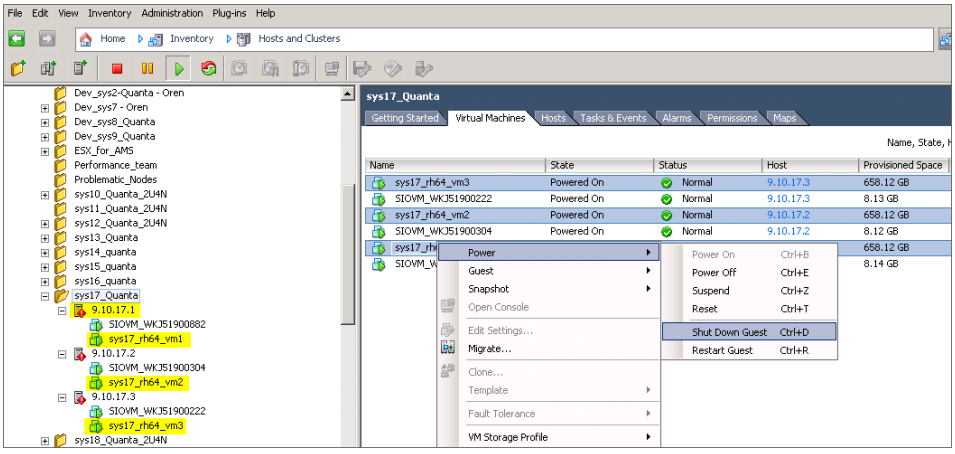

3. The upgrade process will continue, on its own, after the VMs are shut down.

Upgrade fails when NIC in cluster node is down

If any of the NICs in a cluster are down during an upgrade process, the upgrade fails. You must bring the NIC back up before you continue with the upgrade.

Procedure

1. Login to the ESX.
2. Run the CLI command: `esxcli network nic up -n vmnic0`
3. Login to the VxFlex OS GUI and retry the upgrade.

Results

Upgrade successfully completes.

Abort AMS upgrade causes lost connection to SDS

Abort AMS upgrade causes lost connection to SDS and cannot be restored.

Problem: The SDS service stops when the AMS upgrade process is aborted.

Solution:

1. Turn on the ESXi server and Storage VM (SVM).
2. From the SVM run the command: `/opt/emc/scaleio/sds/bin/create_service.sh`.
3. Wait for the re-balance operation to complete.