Dell EMC VxFlex OS

Version 3.x

Troubleshoot and Maintain VxFlex OS

P/N 302-005-638

REV 01

Published April 2019

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

CONTENTS

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

This section contains troubleshooting events and suggested solutions connected to the VxFlex OS system.

Troubleshooting VxFlex OS

This guide contains topics to help you troubleshoot your VxFlex OS system.

This section describes solutions to issues that may arise. Most of the issues requiring troubleshooting are related to events and CLI return messages:

Event notifications
Upon changes or events in the system, VxFlex OS may generate a system event that will be logged in a file. When addressing an event, look for the proper entry in the System Events appendix of the *Monitor VxFlex OS Guide*, according to the name of the event received. For each event that requires attention, the entry will include a description of a possible action to take. Follow the instructions before contacting DELL EMC Support.

CLI messages
When issuing a CLI command, VxFlex OS generates a textual response describing the outcome of the command. In some cases, the result may be a failure. Upon receiving a failure message, you can address it by looking for the proper entry in the Return Messages appendix as described in "Supported alerts and event numbering conventions" of the *Monitor VxFlex OS Guide*, according to the text of the message received.

For each return message that requires attention, the entry will include a description of a possible action to take. Follow the instructions before contacting DELL EMC Support.

Add acceleration pools

You can add Acceleration Pools to a Protection Domain, to accelerate Storage Pool performance.

An Acceleration Pool is a group of acceleration devices within a Protection Domain.

Procedure

1. In the **Backend** > **Devices** view, select the required Protection Domain.
2. Right-click the Protection Domain and select **Add Acceleration Pool**.

   The **Add Acceleration Pool** window is displayed.
3. Enter a name in the **Acceleration Pool Name** box.
4. Select a pool type.
   - For Fine Granularity data layout, select `NVDIMM`. You must have at least one NVDIMM installed in order to select this option.
   - For Medium Granularity data layout, select `SSD`. You must have at least one SSD installed that can be used for the RFcache feature in order to select this option.
5. Click **Add Devices** and then click the **Add device** icon to add a row to the **New Devices** table.
6. Enter the following information in the relevant row of the table:
   - In the **Path** cell, enter the location of the acceleration device
   - In the **Name** cell, enter the name of the acceleration device

- From the **SDS** drop-down list, select the relevant SDS

7. Click **Advanced** and configure the optional settings:
   - **Test and activate device**
   - **Device test timeout: x seconds**
   - **Force devices takeover** (takes over devices that were previously used in a VxFlex OS system.
8. If you want to add more devices, click the **Add device** icon again and configure the fields in the new row.
9. Click **OK**.

Results

The Acceleration Pool has been created, and acceleration devices have been added to it.

After you finish

For RFcache Acceleration Pools, ensure that caching is enabled, using the **Configure Caching** > **Set Read Flash Cache Policy** command. This feature can be enabled at Protection Domain, Storage Pool, or SDS level.

After SDC installation, the Storage VM does not start automatically

Following SDC installation, it is required to set all Storage to start automatically.

After deployment is completed, set all Storage VMs (SVMs) to start automatically:

Procedure

1. Click the ESX **Configuration** tab.
2. From the **Software** section, click **Virtual Machine Startup/Shutdown**.
3. Click **Properties**.
4. In the dialog box, select **Allow virtual machines to start and stop automatically with the system**.
5. Select the SVM and move it to the **Automatic Startup** list.
6. Click **OK**.
7. Repeat this process for all SVMs.

Application server does not see a VxFlex OS volume

Perform the following steps if the application server does not see aVxFlex 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 by running the `scli --query_all_volumes` command.
3. Determine if the SDC is installed on the server:
   - Linux: Run `rpm -qa | grep sdc`
   - ESXi: Run `esxcli software vib list|grep sdc`
   - Windows: Run `Get-Package -Name EMC*`
4. Determine if the SDC is connected to an MDM:

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

5. Ensure that the MDM management IP address is up and running.
6. On an application server, rescan for new volumes:
   - Linux: Run `/opt/emc/scaleio/sdc/bin/drv_cfg --rescan`
   - Windows: Run `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg --rescan`
   - ESXi: Rescan for new devices

Cannot log in to the VxFlex OS Installer after upgrade

Restart the VxFlex OS Gateway service after upgrade to log in.

VMware deployment failures

Correct failures during a 5-node deployment with the VxFlex OS vSphere plug-in deployment wizard.

The deployment may fail with any of the following errors:

- Failed: Install ScaleIO SDS module (Cannot connect to MDM cluster)
- Failed: Install ScaleIO RFCACHE module (Cannot connect to MDM cluster)
- Failed: Configure SDC driver on ESX (Cannot connect to MDM cluster)

To fix these issues, click **Continue deployment** in the wizard. Deployment should continue successfully.

S.M.A.R.T. hardware alerts are not displayed

Hardware alerts are not displayed in the VxFlex OS GUI.

1. Ensure that the proper driver/disk utility is installed on the servers:
   - LSI: storcli
   - HP servers: hpssacli
   - Dell servers: perccli

The CLI utility enables you to display the hardware alerts for the RAID controller and disk status.

2. Ensure that the LIA process is running on the server.
   LIA is used for alert display.
3. Ensure that the specific RAID controller is supported on the server.

VxFlex OS Installer returns an error

VxFlex OS Installer returns an error indicating it may have been disabled.

In the VxFlex OS Installer, if an `HTTP 404 Webpage not available` error (or similar) is displayed, the VxFlex OS Installer may have been disabled.

For more information, refer to “Install and configure the VxFlex OS Gateway in the *Deploy VxFlex OS Guide*.

Installation with the VxFlex OS Installer fails

Error message when installing with the VxFlex OS Installer may indicate that network is inaccessible.

If the following error message appears when installing with the VxFlex OS Installer,

```
Command failed: Could not connect to <IP_address>
```

It could be that this node is not in an accessible network.

Error when unmapping volume from ESXi server

Correct failure that occurs when unmapping volumes from an ESXi server.

Before you begin

You need the login credentials for the ESXi host where the unmapping is not succeeding.

Sometimes, when unmapping volumes from an ESXi server, a " Could not detach LUN" error is displayed. Use the following steps to unmap the volume.

Procedure

1. Use SSH to log in to the ESXi host.
2. Run this command:

   ```
   esxcli storage core device detached remove --all
   ```

3. From VxFlex OS, unmap the volume again.

Results

The unmap command succeeds.

Replacing a faulty accelerated storage device in VxFlex OS

Replace a faulty accelerated (cached) storage device in a VxFlex OS system.

Before you begin

Ensure that you have:

- Administrator authentication credentials to run VxFlex OS SCLI commands
- Access to the VxFlex OS nodes
- Any one of the following parameters:
  - SDS name
  - SDS IP address
  - SDS ID

This procedure applies to VxFlex OS software-only systems. It does not apply to AMS-supported VxFlex OS systems.

Procedure

1. Log in to the scli interface.
2. Identify the failed cached storage HDD device:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Output, similar to the following, appears:

```
Device information (total 7 devices):
         1: Name: N/A  Path: /dev/sdb  Original-path: /dev/sdb
ID: ffcfa55700000000
                Storage Pool: sp1, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         2: Name: N/A  Path: /dev/sdc  Original-path: /dev/sdc
ID: ffcfa55800000001
                Storage Pool: sp1, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         3: Name: N/A  Path: /dev/sdd  Original-path: /dev/sdd
ID: ffcfa55900000002
                Storage Pool: sp1, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         4: Name: N/A  Path: /dev/sde  Original-path: /dev/sde
ID: ffcfa55a00000003
                Storage Pool: sp2, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         5: Name: N/A  Path: /dev/sdf  Original-path: /dev/sdf
ID: ffcfa55b00000004
                Storage Pool: sp2, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         6: Name: N/A  Path: /dev/sdg  Original-path: /dev/sdg
ID: ffcfa55c00000005
                Storage Pool: sp2, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
         7: Name: N/A  Path: /dev/sdh  Original-path: /dev/sdh
ID: ffcfa55d00000006
                Storage Pool: sp3, Capacity: 929 GB Error-
fixes: 0 scanned 0 MB, Compare errors: 0 State: Normal
```

For a faulty device, the command output displays the fault (error) at the end of the output, while the output for a healthy device does not.

3. Remove the faulty device from the Storage Pool:

```
scli --remove_sds_device (--device_id <ID> |((--sds_id <ID> |
--sds_name <NAME> | --sds_ip <IP> [--sds_port <PORT>]) (--
device_name <NAME> | --device_path <PATH>)))
```

4. Physically replace the faulty device, using the relevant system and vendor guidelines.
5. Add a new device, as a cached device, to the Storage Pool:

```
scli --add_sds_device (--sds_id <ID> | --sds_name <NAME> | --
sds_ip <IP> [--sds_port <PORT>])
--device_path <PATH> [--device_name <NAME>] (--
storage_pool_name <NAME>) | --storage_pool_id <ID>)
```

When the device is added to the Storage Pool, the device starts being cached.

6. Verify the status of the replaced device:

```
scli --query_sds --sds_name <NAME> | --sds_ip <IP> |--sds_id
<ID>
```

Example:

```
scli --query_sds --sds_name SDS1
```

The command output should display the device state without errors (that is, the replaced device in healthy state); otherwise, repeat the procedure.

Results

The faulty cached device is replaced in the VxFlex OS system.

VxFlex OS CLI or GUI cannot connect to an MDM

Perform the following steps if the VxFlex OS CLI or GUI cannot connect to an MDM.

Perform the following steps:

Procedure

1. Ping the MDM IP address to ensure you have connectivity.
2. Ensure that you are connecting to the IP address of the Master MDM.

   If the MDM ownership has changed, try to connect to the IP address of the Slave MDM.
3. Check if the MDM is running, by typing the following command:

   ```
   ps -ef| grep mdm
   ```
4. Ensure that the management IP address is up and running.

VxFlex OS Gateway fails to run

The VxFlex OS Gateway fails to run, or it runs but does not listen to any communications. In some cases, gateway installation fails.

Any of the above scenarios may occur due to theVxFlex OS Gateway port collision. The VxFlex OS Gateway requires listening ports 80 and 443 when running as root, or listening ports 8080 and 8443 when running as non-root.

These ports are required in order to receive communications from the OpenStack controller.

On Linux servers, to verify whether the required ports are already in use, run:

```
netstat -tupln
```

On Window servers, run:

```
netstat -nab
```

None of the required ports (80 or 443 for root, or 8080 or 8443 for non-root) should be included in the command output.

If any of the required ports is already in use by an application other than VxFlex OS, the services currently using the ports should be shut down for the duration of the installation and use of the VxFlex OS Gateway. Alternatively, the VxFlex OS Gateway installation should be delayed until the required ports are made available.

The table lists the default ports used by VxFlex OS

SCLI add_sds command fails due to communication error or MDM going offline

When you run the SCLI `add_sds` command, the command may fail either due to a communication error or the Master MDM going offline.

Normally, when you run `add_sds`, the SDS receives a request and returns acknowledgment (ACK) to the MDM. However, after you initiate an `add_sds` command, if the MDM goes down for any reason, the MDM may not receive the ACK message (because the MDM is offline), but the SDS presumes that it is attached to the MDM.

Thus, when the MDM does not receive the ACK message, retrying the `add_sds` command fails with status as `SDS is already attached to this MDM`.

To resolve this issue, it is recommended that you clean the SDS configuration by running the `add_sds` command with `force_clean` flag and then re-run `add_sds`.

Virtual IP feature is not functional

Configuring virtual IP addresses on an MDM cluster via vSphere plug-in may disable the virtual IP feature in the system.

It is recommended that you configure the virtual IP addresses on a physical system only via VxFlex OS Installer. The virtual IP configuration using vSphere plug-in may

lead to communication issues with SDCs, which may trigger IO errors and cause the virtual IP feature to fail.
However, if you have already used the vSphere plug-in instead, you must perform the following steps to resolve the issue:

Procedure

1. Remove the virtual IP addresses by running the following command:

```
scli --modify_cluster_virtual_ips
```

2. Remove the virtual IP interfaces by running the following command:

```
scli --modify_virtual_ip_interfaces
```

3. Edit the IP addresses in the `drv_cfg` file of each SDC in the system.
   - Linux: `/opt/emc/scaleio/sdc/bin/drv_cfg`
   - Windows: `C:\Program Files\emc\scaleio\sdc\bin\drv_cfg`
   - ESXi: Refer to "Update the SDC parameters in VMware based HCI or Compute node".
4. Reconfigure the virtual IP addresses, via IM.

Results

The virtual IP addresses are configured on the MDM cluster in the system, via VxFlex OS Installer, preventing the failure of the virtual IP feature.

The VMware plug-in responds slowly

To solve a slow response of the VMware plug-in, increase memory size and restart the vSphere web client service.

In an environment with a large-scale of volumes, the plug-in response may slow down. To solve this, increase the maximum memory size to 2GB and then restart the vSphere web client service.

- Windows

  In the configuration file, `C:\Program Files\VMware\Infrastructure\vSphereWebClient\server\bin\service\conf\wrapper.conf`, look for a string similar to `= -Xmx` (the line should also start with `wrapper.java.additional`), and change the value to 2048M.

- Linux

  In the configuration file, `usr/lib/vmware-vsphere-client/server/wrapper/conf/wrapper.conf`, look for `wrapper.java.maxmemory=` and change the value to 2048.

  The `wrapper.java.maxmemory` parameter may not exist in vCenter 6.0 `wrapper.conf`, If this is the case, add it manually:

  `wrapper.java.maxmemory=2048`

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 tasks in the *VxFlex OS Configure and Customize Guide*.
- For VxFlex OS software only, see "Analyze the VxFlex OS system using the system analysis tool in the *Deploy VxFlex OS Guide*.

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

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.

Deploying SVM on a node with a management IP address from a different subnet than the node with the SVM template

Due to a VMware limitation, it is not possible to deploy an SVM on a node with a management IP address that is from a different subnet than the node that hosts the SVM template.

When it is necessary to deploy an SVM on a node with a management IP address that is from a different subnet than the node hosting the SVM template, use one of the following workarounds:

- Add the management IP address from the same SVM subnet to the node that hosts the SVM template, and add the new SVM to an existing system.
- Upload the SVM template to the node that will host the new SVM, and add that node to an existing system. Use the template you uploaded to deploy the SVM.

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*.

SSD devices are not recognized in the operating system

In an ESXi Server with a RAID controller, SSD devices are not recognized correctly by the ESXi host and are not displayed as SSDs in the vSphere plug-in.

Procedure

1. Select the appropriate method to identify the devices' type and to correspond it to the device ID.

Option	Description

DirectPath is not configured	Perform the following procedures on the ESXi host:
	a. List the devices and their types:
	`cd /opt/lsi/perccli/`
	`./perccli /c0/eall/sall show`
	The devices are displayed with their DG number. The devices' type is displayed in the **Med** column.
	b. Display the DG number and corresponding VD number:
	`./perccli /c0/vall show`
	c. Run the following command:
	`esxcli storage core path list`
	The VD number is listed in the **Target** column. Identify the row with the value in the **naa** column that matches the Target number.

DirectPath is configured	Perform the following procedures on the SVM:
	a. List the devices and their types:
	`/opt/MegaRAID/perccli/perccli64 /c0/eall/sall show`
	The devices are displayed with their DG number. The devices' type is displayed in the **Med** column.
	b. Display the DG number and corresponding VD number:
	`/opt/MegaRAID/perccli/perccli64 /c0/vall show`
	c. Run the following command:
	`ls -l /dev/disk/by-path/`
	Search for "pci-0000:02:00.0-scsi-0:2:*X*:0", where *X* is the VD number from the previous step.

RAID controller, battery, or disk device hardware errors

In Dell PowerEdge nodes, use the iDRAC System Event Log to look for RAID controller, cache battery, or disk device hardware errors that may lead to data loss or data integrity issues.

Symptoms

The iDRAC System Event Log may display one of four specific hardware errors.

These two errors indicate that the RAID storage controller has failed:

- `The PERC1 battery has failed`
- `Integrated RAID Controller 1 is unable to recover cached data from the Battery Backup Unit (BBU).`

These two errors indicate that one of the disks on the SDS has failed:

- `The rebuild of Disk 0 in Backplane 1 of Integrated RAID Controller 1 failed due to errors on the source physical disk.`
- `A block on Disk 0 in Backplane 1 of Integrated RAID Controller 1 was punctured by the controller.`

Impact

- When either of the two RAID controller errors occurs, it is possible that data in the faulty storage controller was not evicted to the disks.
- When either of the two SDS disk failure errors occurs, data in the RAID RAM cache may not have been evicted to the disks.

For any of the above errors, a special procedure is required during the part replacement in order to avoid data loss or data integrity issues.

Do not power off and power on the server, or perform an server reset, as these actions may contribute to the data integrity issues. Additionally, do not enter the server into SDS maintenance mode, as this may lead to data loss or data integrity issues.

Solution

- If you are replacing a faulted RAID storage controller on a VxFlex Ready Node server, refer to the *VxFlex Ready Node v 2.x Integrated Storage Controller Card Replace FRU Guide* for the specific server model. Ensure that you follow the special procedure for preparing the server for part replacement and returning the node to operation when RAID controller errors occur.
- If you are replacing a faulted disk on a VxFlex Ready Node server, refer to the relevant *VxFlex Ready Node v2.x Disk Replace FRU Guide* for the specific server model and OS type. In order to avoid data loss or data integrity issues, follow any additional instructions for replacing a disk when either of these errors occurs.
- For all other servers, use the replacement procedure for non-VxFlex Ready Node servers, explained below.

Bug with VMware-based products that use the Flash plug-in

There is a bug with VMware-based products that use the Flash plug-in.

Problem:

During the VxFlex OS VMware Installation, after entering valid values in the relevant fields of the **Configure SVM** page, the **Next** button remains grayed out.

Cause: There is an issue related to VMware- based products that use the Flash plug-in.

Solution:

Download a beta version of Flash. This version includes a security fix resolving the vSphere Web Client and other VMware-based products that rely on the Flash plug-in.

Issue with fix deployment configuration when SDC settings not changed

Perform the following steps to solve error after using **Fix deployment configuration** without changing the SDC and/or the virtual IP settings.

Problem:

When running a clean installation after using **Fix deployment configuration**, without changing the SDC and or the virtual IP settings, SDC configuration fails with error:

```
Failed: Configure SDC driver on ESX
```

The installation process can't be completed (unless an environment is configured with no SDCs) and can't be rolled back. If installation is re-run, the same issue occurs.

Solution:

1. Abort installation.
2. Clean the whole environment. For more information, see "Cleaning the VxFlex OS VMware environment and performing a clean install" in the *VxFlex OS Deployment Guide*
3. Re-run the installation.

Delete a certificate from the trust store

Delete a certificate from the trust store

Syntax:

```
keytool -delete -alias [unique_alias] -keystore
[path_to_certificates_folder]/truststore.jks
```

Example:

```
keytool -delete -alias "givenname=mdm, ou=asd, o=emc,
l=hopkinton, st=massachusetts, c=us, cn=centos-6.4-adi5" -
keystore C:\Users\cj\AppData\Roaming\EMC\scaleio\certificates
\truststore.jks
```

Workaround for insufficient storage space for logs on Linux machines running VxFlex OS

Workaround for insufficient storage space for logs on Linux machines running VxFlex OS

Symptoms

Not enough storage space for logs of one or more of the VxFlex OS components on Linux machines running VxFlex OS.

Impact

Components of VxFlex OS including LIA, SDS, MDM and SDR begin to run out of disk space for logs, making it difficult to troubleshoot problems.

The issue may become more apparent when raising the trace level or increasing the number of trace files.

Solution

Perform the following steps on a Linux machine:

1. Copy the logs and cfg sub-directories of the relevant component to a different location. You can copy the directory to a different disk or another partition.
2. From the machine, go to the respective path of the component.

   In this example, <sds> is used: Copy `/opt/emc/scaleio/<sds>/logs/ and /opt/emc/scaleio/<sds>/cfg/` directories to another directory.

   For example, copy the logs and cfg directories to `/tmp/sds/`.
3. Change the `log_dir` and `cfg_dir` parameters in the `/opt/emc/scaleio/<component>/bin/run_bin.sh` file to the new location of the logs and `cfg` directories, respectively.
4. Restart the service.

Workaround for error handling of adding SDS devices during installation of VxFlex OS with DirectPass I/O mode on vCenter Server

Workaround for error handling of adding SDS devices during installation of VxFlex OS with DirectPass I/O mode on vCenter Server

Symptoms

When installingVxFlex OSwith DirectPass I/O mode, during the step to add a device to a single SDS, the `Enter SVM Credentials` dialog appears, requiring you to enter credentials to connect to the SVM. Regardless of entering the correct credentials, the error `Plugin can't connect to SVM` appears which prevents you from adding devices.

Impact

This issue prevents you from adding devices to VxFlex OS with DirectPass I/O mode via the Vsphere plugin.

Solution

There are two possible solutions to workaround this issue:

Solution A:

- Windows OS - Restart the vCenter service on the machine where vCenter is running. After all vCenter services are up, add devices again.
- Linux - Restart the vCenter Server machine. After vCenter is up, add devices again.

Solution B:

1. From the vCenter Server machine, go to:
   - Windows OS - `C:\Users\vsphere-client\AppData\Roaming\VMware\scaleio\serverConfig.cfg`
   - Linux - `/etc/vmware/vsphere-client/vc-packages/scaleio/serverConfig.cfg`

Note

Save a copy of the `serverConfig.cfg` file before editing.

2. In the `serverConfig.cfg` file, for each SVM, remove the line: `<password>6WTlMZQwsYCQ/KHLXpH3Sw==</password>` and save the file.
3. Restart the vCenter Server machine.
4. Add devices. If the `Enter SVM Credentials` dialog appears, enter SVM credentials.

Deployment fails at Add Device phase

Deployment may fail if a storage device was previously partitioned.

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

Procedure

1. Leave the **Deployment** window open.
2. Open SSH to the management IP of the relevant host and use the root user credentials used during the deployment process.
3. In command line, to find the name of the partitioned storage device, type the command `lsblk`

   Check the output for partitioned devices, and make a note of their names. In the example, the partitioned disk is `sdb` and you can also see the disk size.

```
[root@19 ~]# lsblk
NAME   MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda      8:0    0 744.6G  0 disk
sdb      8:16   0   1.1T  0 disk
└─sdb1   8:17   0   1.1T  0 part
sdc      8:32   0   1.1T  0 disk
sdd      8:48   0   1.1T  0 disk
```

4. Use the following `dd` command to delete the partition.

   Note

   Make sure not to delete the system disk.

```
[root@19 ~]# dd if=/dev/zero of=/dev/sdb bs=512 count=1
1+0 records in
1+0 records out
512 bytes (512 B) copied, 0.000149431 s, 3.4 MB/s
[root@19 ~]#
```

   Check with `lsblk`:

```
[root@19 ~]# lsblk
NAME   MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda      8:0    0 744.6G  0 disk
```

```
sdb      8:16   0   1.1T  0 disk
sdc      8:32   0   1.1T  0 disk
sdd      8:48   0   1.1T  0 disk
```

`sdb` doesn't have a partition anymore.

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

Restart Gateway before configuring the Set Alert Password when the lockbox directory is deleted or corrupted

When the lockbox directory is deleted or corrupted, restart the Gateway before configuring the "Set Alert Password" option. The Gateway successfully restarts and the lockbox is re-created.

Unregister VxFlex OS system from SRS

You can unregister the VxFlex OS system from the SRS gateway.

Procedure

1. In the web browser, go to the IP address of your system's VxFlex OS Gateway.
2. Log in to the VxFlex OS Gateway.
3. From the **Maintain** tab, click **System Logs & Analysis** and select **Unregister VxFlex OS system from SRS**

   The **Unregister VxFlex OS system from SRS** dialog box is displayed.
4. Enter the MDM password in the **MDM admin password** box.
5. Enter the VxFlex OS Gateway IP address in the **SRS server URL** box.
6. Click **Unregister from SRS** to remove the system from SRS.

Validate ConnectEMC Dial-Home

Verify that ConnectEMC (SRS) is correctly configured and is communicating with Dell EMC Global Services by performing a dial-home test.

This procedure describes how to quickly verify the connectivity between the SRS-GW and SRS servers, either after installation or after registering to SRS. The SRS-GW is located at the customer's site and has a verified connection to SRS servers as part of its installation.

Procedure

1. While logged in to the SRS-gateway as the gpadmin operating system user, remove the contents of the directory `/opt/connectemc/archive`.
2. Create a test event for ConnectEMC using the following command:

```
touch /opt/connectemc/poll/testerror
```

Results

With up to 10 minutes (depending on your polling interval setting), ConnectEMC (SRS) will generate a file in one of the following directories, depending on the test result:
If the connection was successful, a file is saved in `/opt/connectemc/archive`.

If the service fails, a file is saved in `/opt/connectemc/archive/failed`.

CHAPTER 2

VxFlex OS Upgrade Troubleshooting

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

SDC upgrade fails on Windows nodes

In some very rare cases, when using the VxFlex OS Installer to upgrade from VxFlex OS v1.32.6 to v2.0.1.4 to v3.0, the SDC upgrade fails.

The issue is encountered only when following the exact upgrade path described above, using the VxFlex OS Installer. To solve the issue, use the VxFlex OS Installer to upgrade the other VxFlex OS components from v2.0.1.4 to v3.0, then upgrade the SDCs manually.

Procedure

1. Use the VxFlex OS Installer to upgrade from v1.32.6 to v2.0.1.4, as described.
2. Use the VxFlex OS Installer to upgrade from v2.0.1.4 to v2.6. Do not upload the SDC packages.
3. After the rest of the components are upgraded, upgrade the SDCs manually:
   a. Upgrade the SDC.

      ```
      EMC-ScaleIO-sdc-2.6-<build>.X.msi
      ```

   b. Restart the SDC server.

      The SDC is upgraded.

   c. Repeat these steps for every SDC server.

Upgrade of components via VxFlex OS Installer may freeze

Upgrade process may freeze during upgrade of VxFlex OS components through the VxFlex OS Installer.

If the upgrade process freezes, check the following:

1. Communication issues. Fix any communication errors.
2. Verify that all the installation packages were uploaded successfully.
3. Look for other upgrade-related errors.

After fixing the issue, restart the Gateway service, and restart the upgrade.

Upgrade fails because maintenance mode cannot be exited

The presence of a failed SDS device may prevent the SDS to be exited from maintenance mode. This will stop an upgrade in process.

Before you begin

To continue the upgrade without removing the device:

Procedure

1. From the VxFlex OS Installer in the **Monitor** screen, click **Retry failed**.
2. Select **Force (not recommended)**, then confirm the operation.

Results

The upgrade process continues.

VxFlex OS Gateway server recovery during upgrade

The following procedures describe how to recover a VxFlex OS Gateway when the physical host server fails:

Recovering VxFlex OS Gateway on a new server during upgrade

Recover the VxFlex OS Gateway (GW) on a new server and continue with the process of VxFlex OS upgrade, when the VxFlex OS Gateway server fails during VxFlex OS upgrade.

If the VxFlex OS Gateway server fails during an upgrade, the VxFlex OS Gatewayon VxFlex OS Installer will also go down. You must recover the VxFlex OS Gateway first, and then the VxFlex OS Installer. This procedure will cause the MDM to use the new VxFlex OS Gateway in place of the failed one.
To recover the VxFlex OS Gateway on a new server, perform the following steps:

Procedure

1. Install the VxFlex OS Gateway on a new server (identical to the one that crashed), using any of the following options:
   - Linux OS: Use the VxFlex OS Gateway RPM file.
   - Windows OS: Use the VxFlex OS Gateway MSI file.
2. If the LIA-trusted IPs feature was enabled during installation, you must perform the steps described in #unique_40 before continuing with this procedure.
3. Open the `gatewayUser.properties` file, located at `/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/gatewayUser.properties`, and edit the IM parameters.

Parameter	Description	Action
`upgrade.mdm.data.ips`	Configures MDM data IPs	Enter comma-separated IP addresses for different cluster nodes. • If a single node has multiple IP addresses, then enter plus-sign-separated IPs. • Ensure that the sequence of items in the list matches the order of data, management (mgmt), and role for the same node. • Reserve the empty values for Tie Breaker management (TB mgmt) IP addresses. Example: • `upgrade.mdm.data.ips` = 192.100.1.10,192.100.1.49,192.100.1.9,192.100.1.17,192.100.1.11 • `upgrade.mdm.mgmt.ips` = 10.76.60.41,10.76.60.48,10.76.60.41,,
`upgrade.mdm.mgmt.ips`	Configures MDM management IPs	Enter comma-separated IP addresses for different cluster nodes. • If a single node has multiple IP addresses, then enter plus-sign-separated IP addresses

Parameter	Description	Action
		• Ensure that the sequence of items in the list matches the order of data, management (mgmt), and role for the same node • Reserve the empty values for Tie Breaker management (TB mgmt) IP addresses. Example: • `upgrade.mdm.data.ips` = 192.100.1.10,192.100.1.49,192.100.1.9,192.100.1.17,192.100.1.11 • `upgrade.mdm.mgmt.ips` = 10.76.60.41,10.76.60.48,10.76.60.41,,
`upgrade.mdm.role`	Configures MDM roles	Enter the role of the MDM server; mdm, slave, or tb (Tie Breaker). The line of the MDM role should correlate to the MDM data IP and MDM management IP addresses. Example: `upgrade.mdm.role` = mdm,slave,slave,tb,tb
`upgrade.mdm.version.target`	Configures the target version of MDM	Enter the MDM target version to which you want to upgrade. The first item in the `version.target` must represent the first cluster node. Example: `upgrade.mdm.version.target` = 2.0-828.0,2.0-828.0,2.0-828.0
`upgrade.mdm.version.orig`	Configures the original version of MDMs	Enter the MDM's original version, from which you are upgrading. The first item in the `version.orig` must represent the first cluster node. Example: `upgrade.mdm.version.orig` = 1.32-255.0,1.32-255.0,1.32-255.0
`upgrade.mdm.actor.port`	Configures MDM actor ports	Enter comma-separated MDM actor ports, as an ordered list. Each entry is a relevant option, rollback or upgrade. Option selection relates to the operation, version upgrade or rollback that was being performed before the VxFlex OS Gateway crashed. Example: `upgrade.mdm.status` = upgrade
`upgrade.sds.mm_list`	Removes the SDS from Maintenance mode	Change the status of the SDS, which was in maintenance mode before the GW server crash using the `upgrade.sds.mm_list` parameter. Example: `upgrade.sds.mm_list` = 10.76.60.32,10.76.60.33,10.76.60.34,10.76.60.35,10.76.60.36

Results

The IM parameters in the `gatewayUser.properties` file match with the scenario that existed before the VxFlex OS Gateway server crash.

VxFlex OS Gateway recovery of SNMP configuration

Recover SNMP configuration when the VxFlex OS Gatewayy crashes during VxFlex OS upgrade.

There are two methods of recovering SNMP configuration:

- Recovering SNMP configuration using VxFlex OS Installer
- Manually recovering SNMP configuration

Recovering SNMP configuration using VxFlex OS Installer

You can use the VxFlex OS Installer to recover the SNMP configuration, when the VxFlex OS Gateway crashes during a VxFlex OS upgrade.

Procedure

1. Delete the alert service user from the MDM:
   a. Log in to the MDM with the admin user.

      Example:

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

   b. List all existing users:

      ```
      scli --query_users
      ```

   c. Delete the designated user:

      ```
      scli -delete_user (-user_id <ID> | --username <NAME>)
      ```

      Example:

      ```
      scli --delete_user --username alertservice
      ```

   d. List all existing users, and ensure that the deleted user is not listed:

      ```
      scli --query_users
      ```

2. Use the VxFlex OS Gateway to extend the VxFlex OS system, using an updated CSV file.
3. In the **Installation Configuration** window, select **Set advanced options (optional)**.

   The **Advanced Configuration** section is displayed.
4. In the **Config. Options** section, select **Enable alert service (required for SNMP and ESRS)**.
5. In the **Traps Receiver IP** field, type the traps receiver IP address.

6. Complete the extend operation.

Manually recovering SNMP configuration

Recover SNMP configuration manually, when the VxFlex OS Gateway crashes during VxFlex OS upgrade.

Procedure

1. Create a Lockbox.
2. On the VxFlex OS Gateway, edit the `gatewayUser.properties` file:

   For more information, see the *VxFlex OS User Guide*

   a. Change the value of `features.enable_snmp` to *true*.

   b. Enter IP addresses in `snmp.traps_receiver_ip`.

      VxFlex OS supports up to two comma-separated TRAP receivers.

   c. In `snmp.sampling_frequency`, update the MDM sampling period (optional step).

   d. In `snmp.resend_frequency`, update the time period to re-send the already existing TRAPs (optional step).
3. Restart the VxFlex OS Gateway service.

INDEX