

# Release Notes

## Dell EMC VxFlex OS

Version 3.0.1.1

### VxFlex OS 3.0.1.1 Release Notes

REV 04

March 2020

These release notes contain supplemental information about this release.

• <a href="#">Revision history</a> .....	2
• <a href="#">Product description</a> .....	2
• <a href="#">New and changed features</a> .....	2
• <a href="#">Support deprecation</a> .....	4
• <a href="#">Fixed issues</a> .....	4
• <a href="#">Software packages</a> .....	5
• <a href="#">Install and upgrade</a> .....	7
• <a href="#">Known issues and limitations</a> .....	8
• <a href="#">Operating system known issues and limitations</a> .....	19
• <a href="#">Additional resources</a> .....	20

# Revision history

**Table 1** VxFlex OS 3.0.1.1 release notes revision history

Revision	Date	Description
01	January 15, 2020	Release of VxFlex OS 3.0.1.1
02	February 27, 2020	Document refresh
03	March 19, 2020	Document refresh
04	March 24, 2020	Document refresh

## Product description

Dell EMC VxFlex OS, formerly known as ScaleIO Software, is scale-out block storage software that enables customers to create a scale-out server SAN or hyper-converged infrastructure on x86 server hardware.

VxFlex OS uses local storage devices and turns them into shared block storage that has all the benefits of SAN-but at a fraction of cost and complexity. VxFlex OS will continue to support the existing customer installations and will support an upgrade path from ScaleIO software to the VxFlex OS release.

VxFlex Ready Node is a combination of VxFlex OS and Dell PowerEdge® servers, optimized to run VxFlex OS software, enabling customers to quickly deploy a fully architected, scale out server SAN with heterogeneous hypervisor support. It also includes Automated Management Services (AMS).

The lightweight VxFlex OS software components are installed on the application servers and communicate via a standard LAN to handle the application I/O requests sent to VxFlex OS block volumes. An extremely efficient decentralized block I/O flow, combined with a distributed, sliced volume layout, results in a massively parallel I/O system that can scale up to thousands of nodes.

VxFlex OS is designed and implemented with enterprise-grade resilience. Furthermore, the software features an efficient, distributed, self-healing process that overcomes media and node failures, without requiring administrator involvement.

Dynamic and elastic, VxFlex OS enables administrators to add or remove nodes and capacity on-the-fly. The software immediately responds to the changes, rebalancing the storage distribution and achieving a layout that optimally suits the new configuration.

Because VxFlex OS is hardware agnostic, the software works efficiently with various types of disks, including: magnetic (HDD) and solid-state disks (SSD), flash PCI Express (PCIe) cards, networks, and hosts. VxFlex OS can easily be installed in an existing infrastructure as well as in green field configurations.

## New and changed features

Learn about new and changed features in this version of VxFlex OS.

### New features for VxFlex OS

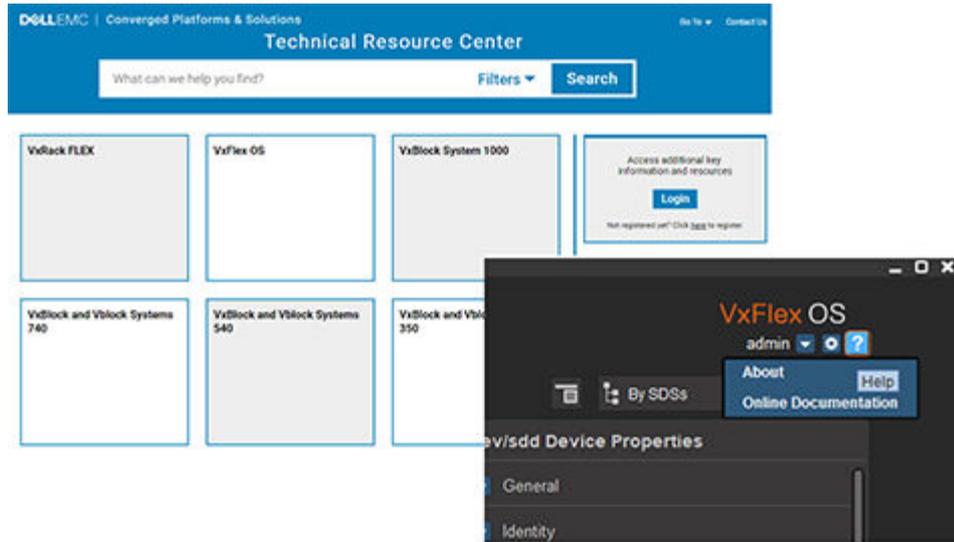
Learn about the new features introduced in VxFlex OS 3.0.1.1 software.

There are no new features included in this release of VxFlex OS.

## New Documentation Set for VxFlex OS v3.0.x, and Online Documentation Portal

The VxFlex OS v3.0 documentation includes a new set of documentation that makes it easier for you to find information and follow common procedures and workflows.

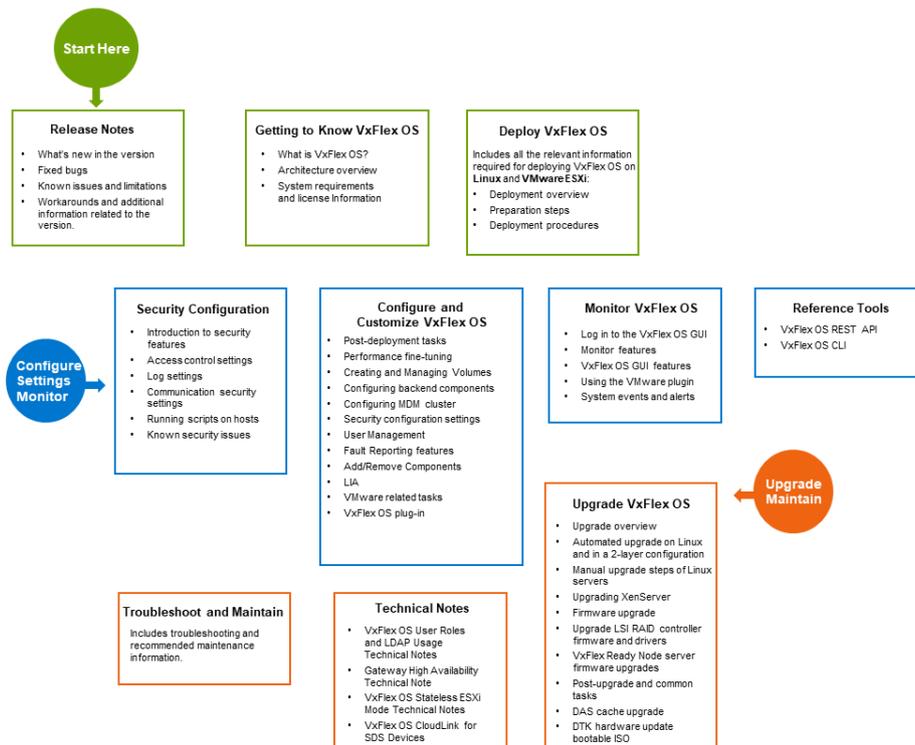
The VxFlex OS v3.0 documentation is now available on the Dell EMC Technical Resource Center; click [here](#) to view and download the latest version of the documentation. You can also access this link from the VxFlex OS GUI, by clicking the ? icon and selecting **Online Documentation**.



For more information on what is included in the new documentation, see the *How to Find Information* document on the Dell EMC Technical Resource Center.

### Welcome to the VxFlex OS Documentation Library

Here you can find information on what is available in the VxFlex OS Documentation. Click [here](#) to view the documentation on the Dell EMC Technical Resource Center.



## Changed features

Learn about the enhanced features for VxFlex OS v3.0.1.1.

The NVDIMM Interleaving setting should be disabled in the BIOS by default (configured in the relevant VxFlex iDM). To verify or set, refer to the following: [https://www.dell.com/support/manuals/us/en/04/poweredge-t640/nvdimn-n\\_ug\\_pub/bios-configuration-settings-for-nvdimn?guid=guid-a69462d3-b86c-4f6b-903e-269f19af591f&lang=en-us](https://www.dell.com/support/manuals/us/en/04/poweredge-t640/nvdimn-n_ug_pub/bios-configuration-settings-for-nvdimn?guid=guid-a69462d3-b86c-4f6b-903e-269f19af591f&lang=en-us)

## Support deprecation

Learn about the features that are no longer supported in VxFlex OS 3.0.1.1.

### AMS compute management support

The Compute and Networking feature is no longer supported on VxFlex Ready Node.

### Windows server backend components

MDM and SDS components are no longer supported on Windows servers (SDC, VxFlex OS Gateway and VxFlex OS GUI are still supported).

### SATADOM boot device support

VxFlex OS v3.0 cannot be deployed on the 32 GB SATADOM boot device that was sold in the first generation of Dell EMC ScaleIO/VxRack Nodes and VxRack Flex (a hardware solution based on Quanta servers). Customers who have older hardware with 32 GB SATADOM but want to use the CentOS SVM in their ESXi configurations on VxFlex OS 2.6.x should contact their account representative to open an RPQ.

### DAS Cache support on SVM

The Storage VM in this release is based on CentOS 7.5, but DAS Cache does not support RHEL\CentOS 7.5. DAS Cache will continue to be supported on Bare Metal RHEL 7.3 and SLES 12.2.

### Multiple SDSs on the same server

The installation of multiple SDS instances on the same server is no longer supported. To guarantee a better customer experience and improved stability, VxFlex OS currently only supports a single SDS per server.

As previous releases of VxFlex OS supported the installation of multiple SDS instances per server (up to four), existing nodes using the multiple-SDS feature should be reconfigured in order to be eligible for an upgrade. Consult with your Dell EMC Professional Services contact to assist with this configuration change, before you upgrade your system.

### Deprecated operating systems

Refer to the ESSM for details.

## Fixed issues

The following table lists the issues that were fixed in VxFlex OS 3.0.1.1. Each table is sorted according to issue severity (from high to low).

 **Note:**

If an issue was reported by customers, the customers' Service Request numbers appear in the "Issue number & SR number" column, and serve to correlate between customer-reported issues and the VxFlex OS (SCI) issue number.

**Table 2** Fixed issues

Issue number & SR Number	Problem summary
SCI-45994 SR# 14805489	In some scenarios, the non-disruptive upgrade does not start successfully due to duplicate LIA IDs.  "Command failed: Two LIAs on different nodes have the same ID. One of the IDs must be change before starting upgrade."
SCI-47080 SR# 15172346	Adding a large amount of SDS components to a Protection Domain might result in a communication error and/or display of the following error: "Could not add SDS on Protection Domain...SDS is already attached to this MDM."
SCI-47791 SR# 15431166	During a Storage Pool rebalance, an MDM role change in a cluster may cause some storage devices to appear as failed. Removing the node and adding it back to the cluster resolves the issue.
SCI-47925 SR# 15547199	Upgrade operation might fail during spare capacity check with an error: "Alert calculated indicates that not enough spare capacity exists to process".
SCI-48002	The VxFlex OS Installer (IM) cannot be used to add devices to a system according to their device ID path (for example, /dev/disk/...).
SCI-48023	When adding many SDS devices at once, the operation may not be successful for all devices.
SCI-48341 SR# 15826685	LIA trace log files might contain sensitive login information.
SCI-49802 SR# 16394618	When a single SDS fails to be added to a cluster during installation, the Installer (IM) does not continue adding the remaining nodes to the cluster.

## Software packages

Learn about the contents of the VxFlex OS 3.0.1.1 software packages.

### Download location

VxFlex OS 3.0.1.1 software packages can be downloaded from the following locations:

VxFlex OS Software:

<https://support.emc.com/products/33925>

ScaleIO Ready Node 13G:

<https://support.emc.com/products/41077>

VxFlex Ready Node 14G:

<https://support.emc.com/products/42216>

VxRack Node:

<https://support.emc.com/products/39045>

### VxFlex OS 3.0.1.1 Linux

- *VxFlex OS 3.0.1.1 COREOS*  
Contains VxFlex OS 3.0.1.1 SDC Core component for COREOS
- *VxFlex OS 3.0.1.1 GPG-RPM-KEY*  
Contains VxFlex OS 3.0.1.1 RPM key (RSA for all Linux flavors and DSA for XEN) used to authenticate the RPM packages (needed for manual installation and upgrade)
- *VxFlex OS 3.0.1.1 Gateway for Linux*  
Contains VxFlex OS 3.0.1.1 Gateway, which includes the following: Installer, Gateway (for REST-API, Alerts, SNMP/ESRS) for Linux/Ubuntu (can be used to deploy VxFlex OS on Linux-based operating systems only). CSV deployment templates are also included.
- *VxFlex OS 3.0.1.1 GUI for Linux*  
Contains VxFlex OS 3.0.1.1 GUI for Linux, which includes DEB and RPM packages, and the software agreement.
- *VxFlex OS 3.0.1.1 RHEL OEL6*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XCACHE) for RHEL6 / CentOS6 / OEL6
- *VxFlex OS 3.0.1.1 RHEL OEL7*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XC) for RHEL7 / CentOS7 / OEL7
- *VxFlex OS 3.0.1.1 RHEL 8*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XC) for RHEL8 / CentOS8 / OEL8
- *VxFlex OS 3.0.1.1 SLES 12.4*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XCACHE) for SLES12 sp4
- *VxFlex OS 3.0.1.1 SLES 12.5*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XCACHE) for SLES12 sp5
- *VxFlex OS 3.0.1.1 SLES 15*  
Contains VxFlex OS 3.0.1.1 Core components for (MDM/SDS/SDC/LIA/XCACHE) SLES 15
- *VxFlex OS 3.0.1.1 Ubuntu 16.04 LTS*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XCACHE) for Ubuntu 16.04
- *VxFlex OS 3.0.1.1 Ubuntu 18.04 LTS*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XCACHE) for Ubuntu 18.04
- *VxFlex OS 3.0.1.1 Citrix Hypervisor (XenServer) 7.1.2 LTSR*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XCACHE) for Citrix Hypervisor (XenServer) 7.1.2 LTSR
- *VxFlex OS 3.0.1.1 Citrix Hypervisor (XenServer) 7.6*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XCACHE) for Citrix Hypervisor (XenServer) 7.6. For XenServer version 7.6, artifacts for XenServer 7.3 should be used. XenServer 7.3 artifacts are compatible with XenServer versions 7.3–7.6.
- *VxFlex OS 3.0.1.1 Citrix Hypervisor (XenServer) 8.0*  
Contains VxFlex OS 3.0.1.1 Core components (MDM/SDS/SDC/LIA/XCACHE) for Citrix Hypervisor (XenServer) 8.0

### VxFlex OS 3.0.1.1 macOS

- *VxFlex OS 3.0.1.1 GUI package for macOS (10.12, 10.13)*  
Contains the VxFlex OS 3.0.1.1 GUI package for macOS.

### VxFlex OS 3.0.1.1 Windows

- *VxFlex OS 3.0.1.1 Complete Windows SW*  
Contains all VxFlex OS 3.0.1.1 Windows-related software artifacts, including Gateway, CSV deployment templates, GUI, and Core components (LIA and SDC)

### VxFlex OS 3.0.1.1 VMware

- *VxFlex OS 3.0.1.1 Complete VMware SW*  
Contains all VxFlex OS 3.0.1.1 VMware-related software artifacts:
  - VxFlex OS 3.0.1.1 ESX  
Contains the OVA file, drv\_cfg file, and SDC package.
  - *VxFlex OS 3.0.1.1 vSphere Plugin*  
Contains VxFlex OS 3.0.1.1 vSphere plugin installer

### VxFlex OS 3.0.1.1 Complete Software Download

Contains all software artifacts for VxFlex OS 3.0.1.1 (same as the above structure)

### VxFlex Ready Node 3.0.1.1 AMS

- *Upgrade .ZIP file*  
Contains the `upgrade.tar` file for upgrading the VxFlex OS Core components (MDM/SDS/SDC/LIA) for ESXi, and Linux, and the `Upgrade-replacesvm-<VERSION>.tar` file for replacing the SVM.
- *VxFlex OS 3.0.1.1 AMS Download Linux .ZIP file*  
Contains the AMS software for Linux and Windows, the AMS GUI, RHEL 7.5 ISO (patched), and the Windows GUI package.
- *VxFlex OS 3.0.1.1 AMS Download ESXi .ZIP file*  
Contains the AMS software for Linux and Windows, the ISO for ESXi 6.7 U2, ESXi 6.5 EP13, and the AMS Windows GUI package.

 **Note:** ISOs are patched with Specter and Meltdown fixes.

- *Standalone packages*  
`upgrade.tar` file
- *Standalone ISOs*  
RHEL 7.3 ISO, patched with Specter and Meltdown fixes.  
RHEL 7.4 ISO, patched with Specter and Meltdown fixes.  
VMware ISO for ESXi 6.0 EP 19  
AMS management ISO (Kylin image), used to deploy a SLES 12.2 image on an R640 management node (AMS). Includes the AMS .RPM install package.

## Install and upgrade

VxFlex OS 3.0.1.1 can be installed on a clean system or used as an upgrade. Learn important information about installing or upgrading this version.

### Installation

For complete instructions for installing VxFlex OS, see the *Deploy VxFlex OS Guide*.

For complete instructions for installing VxFlex Ready Node, see the *VxFlex Ready Node Deployment Guide*.

## Licensing

VxFlex OS use requires a license. Please call your Dell EMC account representative to purchase a license for VxFlex OS.

## Upgrade

The following table lists the versions from which the upgrade can be performed:

**Table 3** Upgrade support matrix

ScaleIO / VxFlex OS base version	Upgrade path	Supporting documentation
Earlier than v2.0.1.4	Consult Customer Support	
v2.0.1.4	<ol style="list-style-type: none"><li>1. Upgrade to VxFlex OS v2.5 (intermediate step).</li><li>2. Upgrade to VxFlex OS v3.0.1.1.</li></ol>	<i>Upgrade VxFlex OS v3.x Guide</i>
v2.5, v2.6.x, v3.0, v3.0.0.x, v3.0.1	Upgrade to VxFlex OS v3.0.1.1	<i>Upgrade VxFlex OS v3.x Guide</i>

To upgrade from earlier versions, contact Dell EMC Support.

The *VxFlex OS Upgrade Guide* and *VxFlex Ready Node Upgrade Guide* contain detailed upgrade instructions for your specific product.

## Supported operating systems

The user documentation contains a list of supported operating systems. For the most current information, see the *Dell EMC Simple Support Matrix* (ESSM) at <https://elabnavigator.emc.com/eln/modernHomeSSM>.

## SDC compatibility

By design, the VxFlex OS SDC in this release is compatible with previously released VxFlex OS/ScaleIO systems (backend) that are currently supported (for example, ScaleIO v2.5 or VxFlex OS v2.6.1.1). In addition, the VxFlex OS system (backend) in this release is backward compatible with previously released SDC versions that are currently supported.

Some limitations may apply to mixed-version systems, such as:

- Fine Granularity Storage Pool-based volumes cannot be mapped to a ScaleIO/VxFlex OS v2.x SDC
- Snapshot Policy management is not available with SDC releases earlier than v3.0.

# Known issues and limitations

The following table lists known issues and limitations that exist in VxFlex OS 3.0.1.1. Each table is sorted according to issue severity (from high to low).

### Note:

If an issue was reported by customers, the customers' Service Request numbers appear in the "Issue number & SR number" column, and serve to correlate between customer-reported issues and the VxFlex OS (SCI) issue number.

**Table 4** Known issues and limitations—AMS

Issue number & SR number	Problem summary	Workaround
SCI-50342	When using iDRAC 3.21.26.22 or lower, AMS upgrade to 3.0.1.1 (iDRAC 4.00.00.00) might fail with the following error message in the AMS and iDRAC Job queue log: 'unable to extract a payload'	Upgrade iDRAC manually to 3.34.34.34.34, and then click Retry in AMS to perform the upgrade to 4.00.00.00. This needs to be done on all relevant nodes. Refer to iDRAC upgrade guidelines on how to perform a manual upgrade of firmware, or refer to the Upgrade VxFlex OS Guide.
SCI-47595	In some scenarios, following a non-disruptive upgrade (NDU) from v2.6.x to v3.0.x, some of the nodes might not exit maintenance mode.	<ol style="list-style-type: none"> <li>1. In AMS, go to the Backend tab -&gt; Storage in the SDS view.</li> <li>2. Look for the SDSs that are in maintenance mode and expand their device lists.</li> <li>3. Find devices that are in error state, right click them and clear device errors.</li> </ol> <p>When all device errors are cleared, the SDSs can exit maintenance mode.</p>
SCI-42253	When there are many requests sent to the 'perccli' utility it may become stuck and respond with empty answer. If this happens during deployment, the deployment may fail during VD creation on the physical disks.	Wait a little until the sampler successfully samples the 'perccli' utility, and then click "Retry".
SCI-39387 SR# 12011318	AMS requires the certificate imported for its "renew certificate" process to be an intermediate CA certificate. However, the AMS can successfully import a non-intermediate-CA certificate, that is one which is not valid for signing further certificates. If that happens, it negatively impacts the ability to manage and support the system when there is need to perform a function which is not available through the AMS or when the AMS is not available.	Generate a new certificate for the AMS that is valid for signing further certificates. Perform the AMS "renew certificate" process using the new certificate.
SCI-41564	In the AMS GUI, when performing an ownership migration, without specifying a Monitor user, AMS will hang in "taking ownership" state. This state blocks the addition of new nodes.	Re-run the query with a Monitor user specified. AMS will return to a normal state after a connectivity error is generated.

**Table 4** Known issues and limitations—AMS (continued)

Issue number & SR number	Problem summary	Workaround
SCI-44806	When performing a "Replace SVM" procedure on a v2.6.X system from SLES 12.2 to CentOS 7.5, the Upload OVA stage might fail after the SVM was successfully replaced. If this happens, the error message "There are no nodes that require OVA deployment" is generated.	Restart the AMS service, and the Replace SVM Upload OVA process should continue as expected.

**Table 5** Known issues and limitations—AMS GUI

Issue number & SR number	Problem summary	Workaround
SCI-16900	When trying to remove a standby MDM, while the SVM is running and the MDM is down, a failure message is returned. This is caused because the AMS is also trying to remove the rpm.	The error message can be ignored. The MDM running on the node is not part of the VxFlex OS system anymore.
SCI-40019	When a new software version is added to the AMS repository it will result in other future operations like "Add Node" to fail with a "version too old" message.	None
SCI-46339	Having a different number of NVDIMMs per node across a cluster may create a problem if one is manually trying to change the default AMS wizard selection of SDS devices at the "Add Devices" stage of the wizard. Manually changing this selection can violate the rules by which AMS assigns SDS devices to Storage Pools. This will create an error alert in the deployment wizard. However, the alerts are not specific enough to indicate how to reverse the last (manual) action.	<ol style="list-style-type: none"> <li>1. Do not click the "Abort" button. If you click "Abort", you will have to start the deployment all over again, including re-installation of the nodes.</li> <li>2. Click "Close".</li> <li>3. In the "Backend &gt; Devices" view, filter the display using "by SDSs".</li> <li>4. For each SDS, check how many Acceleration Pools it has. Usually there will be one NVDIMM per Acceleration Pool.</li> <li>5. From "System Settings", click "Add Nodes". You will be back to the last step of the deployment wizard.</li> <li>6. You can only add SDS devices to a Storage Pool accelerated by an Acceleration Pool which has an assigned NVDIMM on the same node.</li> </ol>
SCI-23351	In AMS, when upgrading the GUI (Windows Server) the installation	Avoid changing the installation path in the upgrade wizard.

**Table 5** Known issues and limitations—AMS GUI (continued)

Issue number & SR number	Problem summary	Workaround
	path option is modifiable, but it should be unavailable during an upgrade.	

**Table 6** Known issues and limitations—Gateway

Issue number & SR number	Problem summary	Workaround
SCI-50867	When running the upgrade process via the VxFlex OS Gateway, a problem may occur when upgrading RHEL based nodes. When the upgrade process tries to check the IP address of a node, the following message may appear for some of the RHEL nodes: "Could not get IPs of <ip address>."	VxFlex OS upgrade still relies on the ifconfig command, which should be installed separately on the host. Once it is installed, click Retry on the VxFlex OS Installer window.
SCI-50386	Upgrade from v3.0 to v3.0.1.1 using the VxFlex OS Installer is not successful.	Manually upgrade the LIA component on all nodes before performing the system upgrade procedure.
SCI-12370	In some scenarios, when using IPv6, Installation Manager (IM) might fail to identify that several IP addresses represent the same physical node. This can result in redundant install/upgrade operations.	None
SCI-20141	The "auto collect logs" feature starts due to an error, and prevents the user from doing anything in the VxFlex OS Installer, until log collection is finished.	<ol style="list-style-type: none"> <li>1. Stop the automatic log collection.</li> <li>2. Disable the feature, so that it will not start again until you finish your task.</li> <li>3. Enable the "auto collect logs" feature again.</li> </ol>
SCI-38801	The replace SVM utility does not address static routes configured in the system.	If you have static routes configured in your SLES SVM, reconfigure them after you run "Replace SVM" process on your newly created Centos SVMs.
SCI-42967	When initiating an upgrade of a ScaleIO/VxFlex OS system (from v2.6 and below to v3.0) with a large number of objects in the system, in some rare scenarios, the VxFlex OS Gateway might fail to complete the	Perform an MDM ownership switch to resolve the issue.

**Table 6** Known issues and limitations—Gateway (continued)

<b>Issue number &amp; SR number</b>	<b>Problem summary</b>	<b>Workaround</b>
	retrieval of the system topology from the MDM.	
SCI-47833	The SVM patching feature might not work properly on a Windows Server-based VxFlex OS Gateway.	Use a Linux-based VxFlex OS Gateway.
SCI-5466	Configuration changes of an existing system using a modified CSV file are not supported.	Use one of the VxFlex OS management user interfaces to configure the system.
SCI-40372	During the VxFlex OS Gateway rpm upgrade flow, when lockbox is configured, the upgrade displays an error message: "missing information in mdm credentials (username or password) - cannot update lockbox".	Ignore this message, as the Lockbox has been configured successfully.
SCI-13157	When trying to collect logs using the VxFlex OS Installation Manager while the system utilizes 100% of the disk space on all nodes, the log collection operation takes a very long time, and eventually a misleading time-out error is returned.	None

**Table 7** Known issues and limitations—GUI

<b>Issue number &amp; SR number</b>	<b>Problem summary</b>	<b>Workaround</b>
SCI-41101	Mapping a very large amount (thousands) of volumes to an SDC in a single GUI session might fail due to connectivity issues.	Use CLI/API and/or map volumes gradually in smaller groups
SCI-43176	During a system upgrade, while trying to log in using the GUI, an "Internal Error #34" message might appear in the login window.	Retry login.
SCI-41817	In the GUI, when all snapshots are expanded in a V-tree, the 60th snapshot cannot be viewed.	Navigate to the required snapshot from this preset and then switch to the Volumes preset in order to see all the relevant information.
SCI-42781	When configuring NVDIMMs in systems using the GUI, all discovered memory modules should be assigned to Acceleration Pools, and cannot be left undefined.	Before adding NVDIMM devices, from the Unmount NVDIMM option, choose a region or regions and unmount only those regions. The rest of the regions will be available for AMS. Note: There is a region per single NVDIMM.

**Table 7** Known issues and limitations—GUI (continued)

Issue number & SR number	Problem summary	Workaround
SCI-42944	VxFlex OS v3.0 GUI connected to a ScaleIO v2.0.1.x cluster is missing the Inflight checksum option	Use the CLI to enable checksum for a storage device.

**Table 8** Known issues and limitations—MDM

Issue number & SR number	Problem summary	Workaround
SCI-47388	<p>RFcache devices are not addressed when invoking the CLI command: <code>--update_sds_original_paths</code>.</p> <p>The command updates path configuration for all devices with the currently assigned SDS devices path.</p>	<p>Use the following SCLI command to fix the device path:</p> <pre>scli --update_device_original_path</pre> <p>Usage: <code>scli --update_device_original_path (--device_id &lt;ID&gt;   ((--sds_id &lt;ID&gt;   --sds_name &lt;NAME&gt;   --sds_ip &lt;IP&gt; [--sds_port &lt;PORT&gt;]) (--device_name &lt;NAME&gt;   --device_path &lt;PATH&gt;)))</code></p> <p>Description: Changes the device's original path configuration to the path currently assigned to the device</p> <p>Parameters:</p> <ul style="list-style-type: none"> <li><code>--sds_id &lt;ID&gt;</code> SDS ID</li> <li><code>--sds_name &lt;NAME&gt;</code> SDS name</li> <li><code>--sds_ip &lt;IP&gt;</code> SDS IP address</li> <li><code>--sds_port &lt;PORT&gt;</code> Port assigned to the SDS</li> <li><code>--device_id &lt;ID&gt;</code> Device ID</li> <li><code>--device_name &lt;NAME&gt;</code> Device Name</li> <li><code>--device_path &lt;PATH&gt;</code> SDS storage device path or file path</li> </ul>
SCI-11969	Under heavy load, on a VxFlex OS installed device on a slave MDM, the slave MDM might become temporarily out-of-sync (degraded). This happens because Master MDM updates cannot be written to the device within a one-second timeout period.	As degradation is temporary, the risk to the system is minimal and is auto-corrected. To avoid this issue, it is recommended to install VxFlexOS MDM on faster media (NVMe/SSD). In larger systems, it is also recommended to install the MDMs on separate machines.
SCI-16315	After configuring virtual IP addresses, if the Master MDM discovers that its virtual IP	Once the network problem is fixed, start the MDM processes again using <code>create_service</code> .

**Table 8** Known issues and limitations—MDM (continued)

Issue number & SR number	Problem summary	Workaround
	addresses are unreachable, it will try to perform a switch-over. Virtual IP addresses may be unreachable because the data network switch is down and the cluster is actually using a different network. If no MDM is able to obtain the virtual IP addresses, the MDM processes might shut down.	
SCI-19588	In an extremely non-uniform Storage Pool configuration (SDS or Fault Set that accounts for more than half of the Storage Pool capacity), some of the capacity will not be used by the system, even though it appears to be "free".	
SCI-41601	The snapshot policy mechanism allows for offsets of up to 1 minute in snapshot creation, due to delays caused by MDM reboots and switch-overs. This might present an inconsistent snapshots view.	The snapshot's creation time is a display-only attribute of the snapshot, and has no affect on snapshot maintenance or the order in which snapshots are eventually deleted
SCI-42408	When using the snapshot capability with Fine Granularity (FG) Storage Pools, the Base volume physical capacity and Snapshot physical capacity size calculation (post compression) in the GUI/CLI might not be accurate.	This issue will be addressed in a future release.
SCI-8508	It is not possible to add an MDM when the network latency is greater than 200 msec.	None
SCI-12999	When an LDAP user is assigned to both Security and BackEndConfig groups, upon CLI login, the message "User role is SuperUser" appears, even if the user is not assigned to all the groups.	Since the actual permissions are set according to the assigned groups, the message can be ignored.
SCI-27564	Original snapshot deletion time might not represent the actual deletion time. Deviations might occur due to MDM crashes or switch-overs of up to 1 minute per crash.	Actual deletion time can be calculated from system information with the help of Customer Support.
SCI-11046	When there are device-related and SDS-related oscillating failures in the system and an MDM switch-over	When an additional MDM switch-over occurs, these oscillating failures counters will be available.

**Table 8** Known issues and limitations—MDM (continued)

Issue number & SR number	Problem summary	Workaround
	occurs, those oscillating failures may not be updated in the current Master MDM.	
SCI-14632	Changing the size of a device that is in use by an SDS is not supported.	To resize a device, first remove the device from the VxFlex OS system, resize it, and add it back to the system.
SCI-21795	By default, the SDS will not add new devices that allow less than ~50 MB/s in 1 MB writes.  During the disk initialization phase, the SDS writes ~200 MB. This might be prominent when using multiple partitions on a device.	When using multiple partitions, add the devices to the SDS, one-by-one. Change the SDS add new device timeout by changing the following parameter in conf.txt: mdm_to_tgt_net__send_timeout (in milli-seconds)

**Table 9** Known issues and limitations—Network

Issue number & SR number	Problem summary	Workaround
SCI-11405	SDS IP addresses must not be ambiguous. For instance, 127.0.0.1 must not be used, as it refers to several machines.	None
SCI-12038	The MDM and SDSs might restart due to a known issue in glibc version 2.12-1.166 or earlier of RH6. The issue is likely to occur when there is heavy traffic on the network.	Update the glibc to version 2.12-1.167. More information can be found in Red Hat Bugzilla (Bug 1243824): <a href="https://bugzilla.redhat.com/show_bug.cgi?id=1243824">https://bugzilla.redhat.com/show_bug.cgi?id=1243824</a>
SCI-27540 SR# 08520639, 07724183	The SDS connectivity test (SDS network test) tool might return inconsistent results in networks with configuration issues (Routing, MTU, etc), and when non-vxFlex OS traffic is running on the data subnet (SDS-SDS, SDC-SDS).	None.

**Table 10** Known issues and limitations—SDC

Issue number & SR number	Problem summary	Workaround
SCI-50039 SR# 16490613	In SLES 12.4 installations, SCSI commands issued via the SG_IO ioctl to /dev/scini* devices will fail and cause a kernel "oops".	None.

**Table 10** Known issues and limitations—SDC (continued)

Issue number & SR number	Problem summary	Workaround
SCI-11026	An ESXi host might not recognize a VxFlex OS volume resize operation.	Perform a re-scan of the ESXi host storage adapters.
SCI-2763	When uninstalling a Linux SDC while I/O is running, the process might fail and generate the following error message: "Module scini is in use".	Reboot the node.

**Table 11** Known issues and limitations—SDS

Issue number & SR number	Problem summary	Workaround
SCI-44997	Addition of pre-partitioned NVMe disks to a VxFlex OS system will cause removal of the partitions instead of failing the operation.	Prior to adding new disks to a VxFlex OS system, make sure that they do not contain any valuable data.
SCI-44515	In rare cases, when deleting a large number of volumes with snapshots while an SDS reboot occurs, the deletions can be finished in the absence of the rebooting SDS. That would follow with the devices in the SDS are automatically attached as "new" devices. Despite being marked as new, these devices still have data residing in NVRAM (from before the reboot). This data can be erased only after the devices finish their attachment as "new" devices. The SDS does not attach the devices because it does not have enough space in NVRAM for both the old NVRAM data and the new data.	Remove the disks from the SDS and add them back again.
SCI-38954	In a hyper-converged Linux environment, if more than 2,000 volumes are mapped to a given SDC, restarting the SDS on the same machine may cause the SDS devices on the machine to enter an error state.	This error state can be resolved by using the "clear device error" command.
SCI-43259	An attempt to migrate a volume towards an unbalanced Storage Pool where one of the devices is completely full will produce a "No space in destination SP" message.	Make sure that the target Storage Pool is balanced before migrating volumes.

**Table 11** Known issues and limitations—SDS (continued)

Issue number & SR number	Problem summary	Workaround
SCI-15736	When almost all capacity in a VxFlex OS system is used, and the system is in maintenance mode, read I/Os may fail. This is due to the fact that in order to assure consistent reads, when a read is performed to a new location, the copy to the temporary copy must be written as well.	None
SCI-44410	Volume snapshot deletion seems stuck or may take long time to complete.	Snapshot deletion is dependent on the system status, and will not complete until system rebuild is over.
SCI-35732	When a disk has failed in an ESXi HCI node, the Storage VM might freeze. This will result in SDS failure, and commencement of a rebuild operation.	<ol style="list-style-type: none"> <li>1. Shut down the SVM.</li> <li>2. Enter the ESXi host into maintenance mode (Shut down or migrate any VM located on the host).</li> <li>3. Reboot the host.</li> <li>4. Identify the faulty device and remove it from the SVM, using "edit virtual machine".</li> <li>5. Start the SVM. The SDS should start, the device should be removed and rebuild should be initiated.</li> </ol>
SCI-3526	Multipath devices cannot be added as SDS devices.	None

**Table 12** Known issues and limitations—vSphere VxFlex OS plug-in

Issue number & SR number	Problem summary	Workaround
SCI-26137	When re-mapping a volume to an SDC on an ESXi node, the device appears to be in detached state. This occurs because when the vSphere VxFlex OS plug-in unmaps the volume, it first detaches the device to make sure that it is not being used, and the ESXi "remembers" that detached state.	Perform an attach command on the device using vSphere web client, PowerCLI, etc.
SCI-15183	The vSphere VxFlex OS plug-in does not allow unmapping a volume from the SDC when the SDC is disconnected.	Unmap using the VxFlex OS GUI or VxFlex OS CLI.
SCI-28108	vSphere VxFlex OS plug-in: When deploying a system with a mix of storage and acceleration devices of	Do not use VMDK mixed environments, because they are not supported.

**Table 12** Known issues and limitations—vSphere VxFlex OS plug-in (continued)

Issue number & SR number	Problem summary	Workaround
	both VMDK and RDM datastores, the deployment is not successful, and generates the error: "Cannot create datastore. Error details: VI SDK invoke exception:com.vmware.vim25.HostConfigFault"	
SCI-38905	When installing VxFlex OS using the vSphere VxFlex OS plug-in, and rolling back from a failed installation, upon re-launching the installation wizard, some of the previously chosen configuration parameters might be missing.	Cancel the operation, and start deployment again.
SCI-9912	In VMware environments, when an MDM cluster configuration fails, only the 'Roll-Back entire deployment' button appears. There is no 'Roll-Back failed Tasks' option.	Roll back the entire system and re-deploy.
SCI-13862	When using the vSphere VxFlex OS plug-in to add SDS devices to an existing system, the plug-in uses the device ID identifier (which is not unique across nested/virtual ESXs) to check if the device was already added. Attempting to add an SDS device post-deployment might fail.	Try to perform this operation using the VxFlex OS GUI or CLI.
SCI-19609	In an ESXi environment, when the SDC was installed manually using the CLI, if you attempt to upgrade the SDC using the vSphere VxFlex OS plug-in, the SDC upgrade fails with a 1009 error message. This indicates that an unexpected error was encountered.	Set a name for the SDC to be used by the VxFlex OS system, in the following format: ESX-<IP_ADDRESS_OF_ESX> For example: ESX-10.103.110.54 Alternatively, upgrade the SDC manually.
SCI-22879	In the vSphere VxFlex OS plug-in, if the password field of a non-selected ESXi is empty in the "Pre-Deployment Actions" screen, the "Run" button is disabled and the operation cannot be started.	Enter the vCenter/datacenter password, and it auto-fills the ESXi hosts below them.
SCI-27158	In the vSphere VxFlex OS plug-in, during a device removal from an Acceleration Pool, it might not be possible to close/cancel the pop-up window. The issue occurs when trying to exit the credentials page.	Click the OK button with the correct credentials, or refresh the web browser.

**Table 12** Known issues and limitations—vSphere VxFlex OS plug-in (continued)

Issue number & SR number	Problem summary	Workaround
SCI-33443	In some cases, when performing a restart to the vSphere-client service, the vSphere VxFlex OS plug-in is deleted from the vCenter.	Register the plug-in again to the vCenter. Refer to the Deploy VxFlex OS Guide for the detailed procedure.
SCI-33572	In some cases, when upgrading the vSphere VxFlex OS plug-in (by unregistering the old version and registering the new version), the plug-in remains in the old version.	The browser's cache needs to be cleared. After clearing the browser's cache, the SWF file will be automatically downloaded again.
SCI-38603	In the vSphere installation wizard, switching back and forth between installation screens during deployment might miss the option of replicate selection (in the Add devices screen).	Close the wizard and restart the deployment.
SCI-26831	Use of the vSphere VxFlex OS plug-in to map a non-named volume to an SDC fails.	Make sure the volume to be mapped has a name prior to mapping it.
SCI-35880	During plugin deployment, in some cases an error is raised because of timeout, and the following message is displayed: "Failed to setSdsPerformanceProfile - SDS does not exist."	Click Retry to continue with the deployment.
SCI-7385	When running the PluginSetup script, a message may appear indicating that the script is not trusted. The script is trusted. It is possible to select Always trust, and this message will not be shown again.	None

## Operating system known issues and limitations

Learn about known issues and limitations for specific operating systems.

**Table 13** Known issues and limitations—operating systems

Issue number & SR number	Problem summary	Workaround
VXFV-157	NVMe disks are not supported on RHEL 8.0 due to the following issue: <a href="https://access.redhat.com/solutions/4280341">https://access.redhat.com/solutions/4280341</a>	None

**Table 13** Known issues and limitations—operating systems (continued)

Issue number & SR number	Problem summary	Workaround
	Failure to observe this limitation might cause RHEL-based SDSs with NVMe to experience kernel panic during a disk removal flow.	
N/A	The appropriate VMware licensing is required for persistent memory support; an Enterprise Plus license might be required.	To check whether your license includes persistent memory support, refer to VMware <a href="#">Compare vSphere Editions and Features</a> .

## Additional resources

Use these resources to find more information about this product, get support and provide feedback.

### 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/>

### Where to get 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.

Copyright © 2019-2020 Dell Inc. or its subsidiaries All rights reserved.

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 Technologies, Dell, EMC, 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.