Symantec™ VirtualStore 6.0.4 Release Notes - Linux

March 2014

Symantec

Symantec™ VirtualStore Release Notes

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 6.0.4

Document version: 6.0.4 Rev 2

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information

- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
  - Error messages and log files
  - Troubleshooting that was performed before contacting Symantec
  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec website.

https://sort.symantec.com/documents

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

http://www.symantec.com/connect/storage-management

Symantec VirtualStore Release Notes

This document includes the following topics:

About this document

This document provides important information about Symantec VirtualStore (SVS) version 6.0.4 for Linux. Review this entire document before you install or upgrade SVS.

The information in the Release Notes supersedes the information provided in the product documents for SVS.

This is "Document version: 6.0.4 Rev 2" of the *Symantec VirtualStore Release Notes*. Before you start, make sure that you are using the latest version of this guide. The latest product documentation is available on the Symantec Web site at:

https://sort.symantec.com/documents

For the latest information on updates, patches, and known issues regarding this release, see the following TechNote on the Symantec Technical Support website:

http://www.symantec.com/docs/TECH141448

Component product release notes

In addition to reading this Release Notes document, review the component product release notes before installing the product.

Product guides are available at the following location on the software media in PDF formats:

`/docs/product_name`

Symantec recommends copying the files to the `/opt/VRTS/docs` directory on your system.

This release includes the following component product release notes:

- *Veritas Storage Foundation Release Notes* (6.0.4)
- *Veritas Cluster Server Release Notes* (6.0.4)
- *Veritas Storage Foundation Cluster File System High Availability Release Notes* (6.0.4)

About Symantec VirtualStore

Symantec VirtualStore (SVS) powered by Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) serves as a highly scalable, highly available NAS solution optimized for deploying and hosting virtual machine. VirtualStore is built on top of Cluster File System (CFS), which provides high availability and linear scalability across the cluster.

Important release information

- For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:
  http://www.symantec.com/docs/TECH164885

- For the latest patches available for this release, go to:
  https://sort.symantec.com/
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit the following URL:
  http://www.symantec.com/docs/TECH170013
  Before installing or upgrading Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

Changes in this release

This section describes the changes introduced in this release.

Support for SLES11 SP3

SVS now supports SUSE Linux Enterprise Server 11 Service Pack 3.

See "Supported Linux operating systems " on page 9.

System requirements

This section describes the system requirements for this release.

Supported Linux operating systems

This section lists the supported operating systems for this release of Veritas products. For current updates, visit the Symantec Operation Readiness Tools Installation and Upgrade page: https://sort.symantec.com/land/install_and_upgrade.

Table 1-1 shows the supported operating systems for this release.

Table 1-1 Supported operating systems

Operating systems	Levels	Kernel version	Chipsets
SUSE Linux Enterprise 11	SP2, SP3	3.0.13-0.27 3.0.76-0.11	64-bit x86, EMT*/Opteron 4.1 64-bit only

* Extended Memory Technology

Note: Only 64-bit operating systems are supported.

If your system is running an older version of SUSE Linux Enterprise Server, upgrade it before attempting to install the Veritas software. Consult the SUSE documentation for more information on upgrading or reinstalling your operating system.

Symantec supports only SUSE distributed kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel Application Binary Interface (ABI) compatibility.

Required Linux RPMs for SVS

Make sure you install the following operating system-specific RPMs on the systems where you want to install or upgrade SVS. SVS will support any updates made to the following RPMs, provided the RPMs maintain the ABI compatibility.

Table 1-2 lists the RPMs that SVS requires for a given Linux operating system.

Table 1-2 Required RPMs

Operating system	Required RPMs
SLES 11 SP2	coreutils-8.12-6.19.1.x86_64.rpm ed-0.2-1001.30.1.x86_64.rpm findutils-4.4.0-38.26.1.x86_64.rpm glibc-2.11.3-17.31.1.x86_64.rpm glibc-32bit-2.11.3-17.31.1.x86_64.rpm ksh-93u-0.6.1.x86_64.rpm libacl-2.2.47-30.34.29.x86_64.rpm libacl-32bit-2.2.47-30.34.29.x86_64.rpm libgcc46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm libgcc46-4.6.1_20110701-0.13.9.x86_64.rpm libncurses5-5.6-90.55.x86_64.rpm libstdc++46-32bit-4.6.1_20110701-0.13.9.x86_64.rpm libstdc++46-4.6.1_20110701-0.13.9.x86_64.rpm module-init-tools-3.11.1-1.21.1.x86_64.rpm pam-32bit-1.1.5-0.10.17.x86_64.rpm parted-2.3-10.21.18.x86_64.rpm

Table 1-2 Required RPMs *(continued)*

Operating system	Required RPMs
SLES 11 SP3	coreutils-8.12-6.25.27.1.x86_64.rpm
	ed-0.2-1001.30.1.x86_64.rpm
	findutils-4.4.0-38.26.1.x86_64.rpm
	glibc-2.11.3-17.54.1.x86_64.rpm
	glibc-32bit-2.11.3-17.54.1.x86_64 .rpm
	ksh-93u-0.18.1.x86_64.rpm
	libacl-2.2.47-30.34.29.x86_64.rpm
	libacl-32bit-2.2.47-30.34.29.x86_64.rpm
	libgcc_s1-32bit-4.7.2_20130108-0.15.45.x86_64.rpm
	libgcc_s1-4.7.2_20130108-0.15.45.x86_64.rpm
	libncurses5-5.6-90.55.x86_64.rpm
	libstdc++6-32bit-4.7.2_20130108-0.15.45.x86_64.rpm
	libstdc++6-4.7.2_20130108-0.15.45.x86_64.rpm
	module-init-tools-3.11.1-1.28.5.x86_64.rpm
	pam-32bit-1.1.5-0.10.17.x86_64.rpm
	parted-2.3-10.38.16.x86_64.rpm

Supported VMware software versions

- VMware vSphere 4 (ESX 4.0 Update 1 and later with vCenter Server 4.0 Update 1 and later)
- VMware vSphere 4.1 (ESX 4.1 and later with vCenter Server 4.1 and later)
- VMware vSphere 5.0 (ESX 5.0 and later with vCenter Server 5.0 and later)
- VMware vSphere 5.1 (ESX 5.1 and later with vCenter Server 5.1 and later)

Supported guest operating system for guest operating system customization while cloning

- Windows XP
- Windows Server 2003
- Windows 7

- Windows Server 2008
- Red Hat Enterprise Linux (RHEL 5)
- Red Hat Enterprise Linux (RHEL 6)
- SUSE Linux Enterprise Server (SLES 10)
- SUSE Linux Enterprise Server (SLES 11)

Note: Customization of some guest operating systems and versions requires vCenter Server to be of sufficient version. Refer to http://www.vmware.com/pdf/vsphere4/r40/vsp_compatibility_matrix.pdf for details.

Supported guest operating systems for VMware View integration while cloning

- Windows XP
- Windows 7

Supported Citrix XenDesktop version

- Citrix XenDesktop 5

Fixed issues

This section includes the issues fixed since the previous major release. The fixed issues are presented in separate tables for each applicable minor release.

Symantec VirtualStore fixed issues

This section describes Symantec VirtualStore issues fixed since the previous major release.

Symantec VirtualStore: issues fixed in 6.0.4

There are no Symantec VirtualStore fixed issues in 6.0.4.

Symantec VirtualStore: issues fixed in 6.0.3

There are no Symantec VirtualStore fixed issues in 6.0.3.

Symantec VirtualStore: issues fixed in 6.0.1

There are no Symantec VirtualStore fixed issues in 6.0.1.

Veritas File System fixed issues

This section describes Veritas File System issues fixed since the previous major release.

Veritas File System: issues fixed in 6.0.4

Table 1-3 describes the incidents that are fixed in Veritas File System (VxFS) in 6.0.4.

Table 1-3 Veritas File System 6.0.4 fixed issues

Incident	Description
3312030	The default quota support on Veritas File System version 6.0.4 is changed to 32 bit.
3270210	On SUSE Linux Enterprise Server 11 (SLES 11) SP3, no support is provided for `VRTSfsadv`, `VRTSfssdk`, `VRTSvxfs` and `VRTSvxvs`.
3270200	Support for SLES 11 SP3 is added.
3268931	Support for SLES 11 SP3 is added.
3268916	Support for SLES 11 SP3 is added.
3268908	Support for SLES 11 SP3 is added.
3257467	Support for SLES 11 SP3 is added.
3248955	The system fails to build modules on SLES 11 SP3 for Veritas Oracle Disk Manager (ODM), because it cannot find the `utsrelease.h` file needed to build the ODM module for SP3 on SLES 11.
3247752	The system panics with the following message during the internal stress tests: `Unable to handle kernel paging request at <addr>`
3214816	With the DELICACHE feature enabled, frequent creation and deletion of the inodes of a user may result in corruption of the user quota file.
3149174	Veritas Oracle Disk Manager (ODM) clone shutdown fails with the `ORA-03113: end-of-file on communication channel` error.
3142045	With Oracle 12c version, the Veritas Oracle Disk Manager (ODM) library causes a version mismatch on the RHEL6 platform.

Table 1-3 Veritas File System 6.0.4 fixed issues *(continued)*

Incident	Description
3140990	Requirement for the ability to turn off VxFS's invalidation of pages for some Network File System (NFS) workloads.
3121933	The `pwrite(2)` function fails with the EOPNOTSUPP error.
3101418	The current time returned by the operating system (Oracle error code `ORA-01513`) during Oracle startup is invalid.
3089210	The message `V-2-17: vx_iread_1 filesystem file system inode inode number marked bad incore` is displayed in the system log.
3079215	Oracle RAC Database creation fails with the `Ora-00600 [ksfd_odmio1]` error when Veritas ODM links.
3068902	In case of stale NFS mounts, the `statfs()` function calls on non-VxFS file systems may cause df commands to hang.
3042485	During internal stress testing, the `f:vx_purge_nattr:1` assert fails.
2999493	The file system check validation fails after a successful full fsck during the internal testing with the message: `run_fsck : First full fsck pass failed, exiting`
2991880	In low memory conditions on a VxFS, certain file system activities may seem to be non-responsive.
2983248	The `vxrepquota`(1M) command dumps core.
2977828	The file system is marked bad after an inode table overflow error occurs.
2966277	Systems with high file system activity like read/write/open/lookup may panic the system.
2963763	When the `thin_friendly_alloc()` and `deliache_enable()` functionality is enabled, VxFS may enter a deadlock.
2926684	In rare cases, the system may panic while performing a logged write.
2908391	It takes a long time to remove checkpoints from the VxFS file system, when there are a large number of files present.
2907930	VxFS and VxVM components fail to install post SLES11 SP2 GA kernel versions with 4 digits.
2907923	VxFS and VxVM components fail to install post SLES11 SP2 GA kernel versions with 4 digits.

Table 1-3 Veritas File System 6.0.4 fixed issues *(continued)*

Incident	Description
2907919	VxFS and VxVM components fail to install post SLES11 SP2 GA kernel versions with 4 digits.
2907908	VxFS and VxVM components fail to install post SLES11 SP2 GA kernel versions with 4 digits.
2885592	The `vxdump` operation is aborted on a file system which is compressed using the `vxcompress` command.
2839871	On a system with DELICACHE enabled, several file system operations may hang.
2834192	You are unable to mount the file system after the full `fsck`(1M) utility is run.
2756779	The code is modified to improve the fix for the read and write performance concerns on Cluster File System (CFS) when it runs applications that rely on the POSIX file-record using the fcntl lock.

Veritas File System: issues fixed in 6.0.3

Table 1-4 describes the incidents that are fixed in Veritas File System in 6.0.3.

Table 1-4 Veritas File System 6.0.3 fixed issues

Incident	Description
3004466	Installation of 5.1SP1RP3 fails on RHEL 6.3.
2895743	Accessing named attributes for some files seems to be slow.
2893551	When nfs connections are under load , file attribute information is replaced by question marks.
2881211	File ACLs not preserved in checkpoints properly if file has hardlink.
2858683	Reserve extent attributes changed after vxrestore, only for files greater than 8192bytes.
2857751	The internal testing hits the assert "f:vx_cbdnlc_enter:1a".
2857629	File system corruption can occur requiring a full fsck of the system.
2821152	Internal Stress test hit an assert “f:vx_dio_physio:4,1” on locally mounter file system.
2806466	fsadm -R resulting in panic at LVM layer due to vx_ts.ts_length set to 2GB.

Table 1-4 Veritas File System 6.0.3 fixed issues *(continued)*

Incident	Description
2773383	Read/Write operation on a memory mapped files seems to be hung.
2756779	Write and read performance concerns on CFS when running applications that rely on posix file-record locking (fcntl).
2641438	Modifications to user name space extended attributes are lost after a system reboot.
2624262	fsdedup.bin hit oops at vx_bc_do_brelse.
2611279	Filesystem with shared extents may panic.
2417858	VxFS quotas do not support 64 bit limits.

Veritas File System: issues fixed in 6.0.1

This section describes the incidents that are fixed in Veritas File System in this release.

Table 1-5 Veritas File System fixed issues

Incident	Description
2764861	Uncompress by vxcompress ignores quota limitation.
2753944	The file creation threads can hang.
2735912	The performance of tier relocation using fsppadm enforce is poor when moving a large amount of files.
2712392	Threads hung in VxFS.
2709869	System panic with redzone violation when vx_free() tried to free fiostat.
2682550	Access a VxFS file system via NFS could cause system panic on Linux while unmount is in progress.
2674639	The cp(1) command with the –p option may fail on a file system whose File Change Log (FCL) feature is enabled. The following error messages are displayed: cp: setting permissions for 'file_name': Input/output error cp: preserving permissions for 'file_name': No data available.
2670022	Duplicate file names can be seen in a directory.
2655788	Using cross-platform data sharing to convert a file system that has more than 32k nlinks does not update the vx_maxlink and maxlink_enable tunables.

Table 1-5 Veritas File System fixed issues *(continued)*

Incident	Description
2651922	ls -l command on local VxFS file system is running slow and high CPU usage is seen.
2597347	fsck should not coredump when only one of the device record has been corrupted and the replica is intact.
2584531	vxfs hangs on ls, du and find.
2566875	The write(2) operation exceeding the quota limit fails with an EDQUOT error (Disc quota exceeded) before the user quota limit is reached.
2559450	Command fsck_vxfs(1m) may core-dump with SEGV_ACCERR error.
2536130	fscdsconv fails to convert FS between specific platforms if FCL is enabled.
2272072	GAB panics the box because VCS engine HAD did not respond. The lobolt wraps around.
2086902	Spinlock held too long on vxfs spinlock, and there is high contention for it.
1529708	Formatting issue with the output of vxrepquota.

Installation and upgrades fixed issues

This section describes the installation and upgrade issues fixed since the previous major release.

Installation and upgrades: issues fixed in 6.0.3

This section describes the installation and upgrade issues fixed in 6.0.3.

Table 1-6 Installation and upgrades 6.0.3 fixed issues

Incident	Description
2967125	Eval injection vulnerability in the Digest module before 1.17 for Perl allows context-dependent attackers to execute arbitrary commands via the new constructor.

Installation and upgrades: issues fixed in 6.0.1

This section describes the incidents that are fixed related to installation and upgrades in this release.

Table 1-7 Fixed issues related to installation and upgrades

Incident	Description
2329580	Unable to stop some SFCFSHA processes.
2873102	Perl module error on completion of SFHA installation
2627076	Incorrect server names sometimes display if there is a clock synchronization issue.
2622987	sfmh discovery issue when you upgrade your Veritas product to 6.0.1
2585899	On RHEL, unable to create storage for OCR and Vote disk when using FQDN instead of using only the node name.
2526709	DMP-OSN tunable value not get persistence after upgrade from 5.1SP1 to 6.0.
2088827	During product migration the installer overestimates disk space use.

Known issues

This section covers the known issues in this release.

Symantec VirtualStore issues

The Virtual machine's local Administrator password may be set to blank (2676078, 2676079)

When clones are made of a Windows 2008 Virtual Machine and Guest OS Customization is enabled, the Virtual Machine's local Administrator password is set to blank.

Workaround: There is no workaround for this issue.

The cluster node may panic (2524087)

On SLES 10 SP4, the cluster node may panic while iSCSI initiator access the LUN from the target.

Workaround

There is no workaround at this time.

CFS commands might hang when run by non-root (2403263)

The CFS commands might hang when run by non-root.

Workaround

To resolve this issue

- Use `halogin` command to save the authentication information before running any CFS commands on a non-root sessions.

  When you run the `halogin` command, VCS stores encrypted authentication information in the user's home directory.

NFS resource might not come online while configuring CNFS share (2488685)

If SELinux is configured as `enforcing` or `permissive`, NFS resource might not come online and go into `FAULTED` state while configuring CNFS share `cfsnfssg` service group.

Sample output:

```
# hastatus -sum

-- SYSTEM STATE
-- System               State                Frozen

A  swlx14               RUNNING              0

-- GROUP STATE
-- Group           System   Probed     AutoDisabled    State

B  cfsnfssg        swlx14   Y          N               OFFLINE|FAULTED
B  cfsnfssg_dummy  swlx14   Y          N               OFFLINE
B  cvm             swlx14   Y          N               ONLINE
B  vip1            swlx14   Y          N               OFFLINE

-- RESOURCES FAILED
-- Group           Type                 Resource             System

D  cfsnfssg        NFS                  nfs                  swlx14
```

Workaround

To resolve this issue you need to add the Ethernet port into the trusted list for SELinux.

- In the System Setup->Firewall configuration, select customize.
- In the Trusted device, select the Ethernet port.

VirtualStore machine clones created while the VirtualStore cluster reboots will probably not start (2164664)

In some cases when you clone while rebooting the SVS nodes, you may receive several of the following error messages:

```
clone vms could not start X server
```

Workaround

Delete all the clones that got created while the node crashed and redo the cloning operation.

Cloning may not work (2348628)

If you cannot clone and you are using the VMware vAPP and OVF templates, then you must disable the vApp.

Workaround

To disable the vAPP

1. In VI Client, right-click on the virtual machine > **Edit Settings** > **Options** > **vApp Options**.
2. Click **Disable**.

Need intelligent NDMP/NBU backups for virtual machines (2378396)

When using NDMP or the NBU client to backup a virtual machine, the space consumed by the backup is equivalent to the size of the disks in the virtual machine, even though not all of the disk space in the virtual machine is used.

If a VMDK (Virtual Machine Disk) file is 10GB in size, but only consumes 1GB of disk space, an backup done by NDMP or the NBU client generates 10GB of backup data, even though the original VMDK file contains 9GB of unassigned disk space.

Workaround

Use VMware-specific backup applications (such as NetBackup for VMware) to create space-efficient backups.

The Symantec Quick Clone Virtual Machine Wizard may not behave as expected when multiple instances are open (2309702)

The wizard may not behave as expected, if you invoke multiple parallel session of the wizard from a single vSphere Client at the same time.

For example, if you do the following:

- Right-click wingoldvm1 and invoke the wizard.
- Then soon after, right-click slesgoldvm1 and invoke the wizard.

This causes you to have two instances of the wizard running from the same vSphere Client and can cause unexpected behavior.

Workaround

To resolve this issue:

- Close both instances of the wizard.
- Reopen a new instance of the wizard.

Virtual machines created by the Symantec Quick Clone Virtual Machine Wizard might not boot correctly if during the process the FileStore cluster node, the ESX Server, or the vCenter Server reboots (2164664, 2374229)

In some cases when you clone using the wizard, and one of the following servers crashes or reboots while the clone process is in progress, the clones might not get created correctly:

- FileStore nodes
- ESX host on which the clones are being created
- vCenter Server

Even if the clones appear in the vCenter inventory as created, the clones GuestOS might not be able to boot.

Workaround

Delete all of the clones that were created when the servers crashed or were rebooted, and redo the wizard operation.

Error message does not always display when you select an incorrect cluster to clone (2372713)

In cases where multiple FileStore clusters are registered with the same Virtual Center, the Symantec Quick Clone Virtual Machine Wizard might not provide a

warning that you selected an incorrect cluster to clone a golden image. This could happen if all of the FileStore clusters are exporting the same file system path, such as `/mnt`. Instead of an advanced warning that you selected the wrong cluster, you instead see an error on the final page of the wizard when the wizard attempts to clone the disks (vmdks) of the golden image. The error that displays is similar to the following example:

```
/mnt/goldvm/goldvm.vmdk no such file or directory...
```

Workaround

There is no workaround for this issue.

The installer output states, "Registering SVS license," even if you enabled keyless licensing

When installing, if you enable keyless licensing, the installer's output includes the following message:

```
Registering SVS license
```

Workaround: This message is harmless and can be ignored. The product will successfully install without a license key.

Veritas File System known issues

This section describes the known issues in this release of Veritas File System (VxFS).

A mutex contention in vx_worklist_lk() can use up to 100% of a single CPU (2086902)

A mutex contention in the `vx_worklist_lk`() call can use up to 100% of a single CPU.

Workaround:

There is no workaround for this issue.

Upgrading from disk layout Version 8 to 9 on a file system with partitioned directories and Storage Checkpoints can return with a read-only file system error message (2583197)

Upgrading from disk layout Version 8 to 9 on a file system with partitioned directories and Storage Checkpoints can return with a read-only file system error message. The issue with partitioned directories occurs because disk layout Version 9 has a

new hash function. The issue with Storage Checkpoints occurs because the Storage Checkpoints are marked as read-only during the upgrade.

Workaround:

Before upgrading a VxFS file system with disk layout Version 8 to Version 9, use the following procedure to avoid this error message.

To avoid the system error message

1. Disable the partitioned directories feature if the feature is enabled by setting the `pdir_enable` tunable to 0.

   See the `vxtunefs`(1M) manual page.

2. Remove all Storage Checkpoints before the upgrade.

   See the `fsckptadm`(1M) manual page.

Using cross-platform data sharing to convert a file system that has more than 32k nlinks does not update the vx_maxlink and maxlink_enable tunables (2655788)

If you use cross-platform data sharing to convert a file system that has more than 32k nlinks, the conversion process does not update the `vx_maxlink` and `maxlink_enable` tunables on the target file system.

Workaround:

After the cross-platform data sharing conversion completes, validate the values of the `vx_maxlink` and `maxlink_enable` tunables. If the file system had more than 32k nlinks before the conversion, ensure that the these tunables are updated on the2437138 target file system before mounting the file system.

Expanding a 100% full file system can cause a panic (2599590)

Expanding a 100% full file system can cause a panic with the following stack trace:

```
bad_kern_reference()
$cold_vfault()
vm_hndlr()
bubbledown()
vx_logflush()
vx_log_sync1()
vx_log_sync()
vx_worklist_thread()
kthread_daemon_startup()
```

Workaround:

There is no workaround for this issue.

Possible assertion failure in vx_freeze_block_threads_all() (2244932)

There is a possible assertion failure in the `vx_freeze_block_threads_all`() call when the `pdir_threshold` tunable is set to 1.

Workaround:

There is no workaround for this issue.

Severe impact in read performance (sequential and random) on compressed files compared to uncompressed files (2609152)

The read throughput is highly degraded for compressed files. The difference is seen for sequential I/O and random I/O. For sequential reads, the degrdataion is visbile even when the amount of data read compressed files is one-third of the uncompressed files (compression ratio).

Workaround:

There is no workaround for this issue.

Taking a FileSnap over NFS multiple times with the same target name can result in the 'File exists' error (2353352)

The "File exists" error occurs as a result of the caching behavior of the NFS client. Because the link operation is successful, the NFS client assumes that a file with the specified target name, such as `file2::snap:vxfs:`, was created.. As a result, the NFS client caches a file with this name.

Workaround: Remove the target file after a snapshot is created. This forces the NFS client to remove the name from the cache. For example:

```
# ln file1 file2::snap:vxfs:
# rm file2::snap:vxfs:
```

Enabling delayed allocation on a small file system sometimes disables the file system (2389318)

When you enable delayed allocation on a small file system, such as around 100 MB, the file system can get disabled. In this case, the following error message ,displays in the system console log:

```
mesg 001: V-2-1: vx_nospace - file_system file system full
(size block extent)
```

Workaround:

Use the `vxtunefs` command to turn off delayed allocation for the file system.

Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system nears 100% usage even if other volumes have free space (2438368)

Delayed allocation sometimes gets turned off automatically when one of the volumes in a multi-volume file system is nearing 100% usage even if other volumes in the file system have free space.

Workaround:

After sufficient space is freed from the volume, delayed allocation automatically resumes.

Task blocked messages display in the console for RHEL6 (2560357)

On RHEL6, the kernel occasionally displays messages in the console similar to the following example:

```
INFO: task seq:16957 blocked for more than 120 seconds.
```

These messages display because the task is blocked for a long time on sleep locks. However, the task is not hung and the messages can be safely ignored.

Workaround: You can disable these messages by using the following command:

```
# echo 0 > /proc/sys/kernel/hung_task_timeout_secs
```

Deduplication can fail with error 110 (2591473)

In some cases, data deduplication fails with a message similar to the following example:

```
Saving    Status    Node            Type        Filesystem
-------------------------------------------------------------------------
00%       FAILED    node01          MANUAL      /data/fs1
        2011/10/26 01:38:58 End full scan with error
```

In addition, the deduplication log contains an error similar to the following example:

```
2011/10/26 01:35:09 DEDUP_ERROR AddBlock failed. Error = 110
```

These errors indicate that the deduplication process is running low on space and needs more free space to complete.

Workaround:

Make more space available on the file system.

vxresize fails while shrinking a file system with the "blocks are currently in use" error (2437138)

The `vxresize` shrink operation may fail when active I/Os are in progress on the file system and the file system is being shrunk to a size closer to its current usage. You see a message similar to the following example:

```
UX:vxfs fsadm: ERROR: V-3-20343: cannot shrink /dev/vx/rdsk/dg1/vol1 -
blocks are currently in use.
VxVM vxresize ERROR V-5-1-7514 Problem running fsadm command for volume
vol1, in diskgroup dg1
```

Workaround:

Rerun the shrink operation after stopping the I/Os.

fsppadm operations issued on multi-volume file system fail if there are other mounted file systems with a disk layout Version less than 6 (2909206, 2909203)

The `fsppadm` command checks all mounted file systems, and if it finds any file systems with a disk layout Version that is less than 6, then it exits with the following error message:

```
# fsppadm assign /dst_vset /tmp/pol_test.xml

UX:vxfs fsppadm: ERROR: V-3-26510: Low level Volume enumeration failure
on / with message  Function not implemented
```

This error occurs because the `fsppadm` command functionality is not supported on a disk layout Version that is less than 6.

Workaround:

There is no workaround for this issue.

LLT known issues

This section covers the known issues related to LLT in this release.

LLT connections are not formed when a vlan is configured on a NIC (2484856)

LLT connections are not formed when a vlan is configured on a NIC that is already used to configure an LLT link.

Workaround: Do not specify the MAC address of a NIC in the `lltab` file while configuring LLT if you want to configure a vlan later. If you have already specified the MAC address of a NIC, then delete the MAC address from the `lltab` file, and update the file before you restart LLT.

LLT may fail to detect when bonded NICs come up (2604437)

When LLT is configured over a bonded NIC and that bonded NIC is DOWN with the `ifconfig` command, LLT marks the corresponding link down. When the bonded NIC is UP again using the `ifconfig` command, LLT fails to detect this change and it doesn't mark the link UP.

Workaround: Close all the ports and restart LLT, then open the ports again.

Cannot use CPI response files to add nodes to a cluster that is using LLT over UDP (2869763)

When you run the `addnode -responsefile` command, if the cluster is using LLT over UDP, then the `/etc/lltab` file generated on new nodes is not correct. So, the procedure fails and you cannot add nodes to a cluster using CPI response files.

Workaround: None

GAB known issues

This section covers the known issues related to GAB in this release.

While deinitializing GAB client, "gabdebug -R GabTestDriver" command logs refcount value 2 (2536373)

After you unregister the port with `-nodeinit` option, the `gabconfig -C` command shows refcount as 1. But when forceful `deinit` option (`gabdebug -R GabTestDriver`) is run to deinitialize GAB client, then a message similar to the following is logged.

```
GAB INFO V-15-1-20239
Client GabTestDriver with refcount 2 forcibly deinited on user request
```

The `refcount` value is incremented by 1 internally. However, the refcount value is shown as 2 which conflicts with the `gabconfig -C` command output.

Workaround: There is no workaround for this issue.

I/O fencing known issues

This section covers the known issues related to I/O fencing in this release.

Fencing may show the RFSM state as replaying for some nodes in the cluster (2555191)

Fencing based on coordination point clients in Campus cluster environment may show the RFSM state as replaying for some nodes in the cluster.

Workaround:

Restart fencing on the node that shows RFSM state as replaying.

After you run the vxfenswap utility the CoordPoint agent may fault (3462738)

After you run the `vxfenswap` utility, if the value of the `FaultTolerance` attribute of the CoordPoint agent is more than the majority (more than 50%) of the coordination points then the Coordination Point agent faults.

Workaround: Manually set the value of the `FaultTolerance` attribute of CoordPoint agent to be less than the majority (more than 50%) of the coordination points.

In absence of cluster details in CP server, VxFEN fails with pre-existing split-brain message (2433060)

When you start server-based I/O fencing, the node may not join the cluster and prints error messages in logs similar to the following:

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1043
Detected a preexisting split brain. Unable to join cluster.
```

In the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
operation failed.
CPS ERROR V-97-1400-446 Un-authorized user cpsclient@sys1,
domaintype vx; not allowing action
```

The `vxfend` daemon on the application cluster queries the coordination point server (CP server) to check if the cluster members as seen in the GAB membership are registered with the CP server. If the application cluster fails to contact the CP server

due to some reason, then fencing cannot determine the registrations on the CP server and conservatively assumes a pre-existing split-brain.

Workaround: Before you attempt to start VxFEN on the application cluster, ensure that the cluster details such as cluster name, UUID, nodes, and privileges are added to the CP server.

The vxfenswap utility does not detect failure of coordination points validation due to an RSH limitation (2531561)

The `vxfenswap` utility runs the `vxfenconfig -o modify` command over RSH or SSH on each cluster node for validation of coordination points. If you run the `vxfenswap` command using RSH (with the `-n` option), then RSH does not detect the failure of validation of coordination points on a node. From this point, `vxfenswap` proceeds as if the validation was successful on all the nodes. But, it fails at a later stage when it tries to commit the new coordination points to the VxFEN driver. After the failure, it rolls back the entire operation, and exits cleanly with a non-zero error code. If you run `vxfenswap` using SSH (without the `-n` option), then SSH detects the failure of validation of coordination of points correctly and rolls back the entire operation immediately.

Workaround: Use the `vxfenswap` utility with SSH (without the `-n` option).

Fencing does not come up on one of the nodes after a reboot (2573599)

If VxFEN unconfiguration has not finished its processing in the kernel and in the meantime if you attempt to start VxFEN, you may see the following error in the `/var/VRTSvcs/log/vxfen/vxfen.log` file:

```
VXFEN vxfenconfig ERROR V-11-2-1007 Vxfen already configured
```

However, the output of the `gabconfig -a` command does not list port b. The `vxfenadm -d` command displays the following error:

```
VXFEN vxfenadm ERROR V-11-2-1115 Local node is not a member of cluster!
```

Workaround: Start VxFEN again after some time.

CP server repetitively logs unavailable IP addresses (2530864)

If coordination point server (CP server) fails to listen on any of the IP addresses that are mentioned in the `vxcps.conf` file or that are dynamically added using the command line, then CP server logs an error at regular intervals to indicate the failure. The logging continues until the IP address is bound to successfully.

```
CPS ERROR V-97-51-103 Could not create socket for host
10.209.79.60 on port 14250
CPS ERROR V-97-1400-791 Coordination point server could not
open listening port = [10.209.79.60]:14250
Check if port is already in use.
```

Workaround: Remove the offending IP address from the listening IP addresses list using the `rm_port` action of the `cpsadm` command.

See the *Symantec VirtualStore Administrator's Guide* for more details.

Fencing port b is visible for few seconds even if cluster nodes have not registered with CP server (2415619)

Even if the cluster nodes have no registration on the CP server and if you provide coordination point server (CP server) information in the `vxfenmode` file of the cluster nodes, and then start fencing, the fencing port b is visible for a few seconds and then disappears.

Workaround: Manually add the cluster information to the CP server to resolve this issue. Alternatively, you can use installer as the installer adds cluster information to the CP server during configuration.

The cpsadm command fails if LLT is not configured on the application cluster (2583685)

The `cpsadm` command fails to communicate with the coordination point server (CP server) if LLT is not configured on the application cluster node where you run the `cpsadm` command. You may see errors similar to the following:

```
# cpsadm -s 10.209.125.200 -a ping_cps
CPS ERROR V-97-1400-729 Please ensure a valid nodeid using
environment variable
CPS_NODEID
CPS ERROR V-97-1400-777 Client unable to communicate with CPS.
```

However, if you run the `cpsadm` command on the CP server, this issue does not arise even if LLT is not configured on the node that hosts CP server. The `cpsadm` command on the CP server node always assumes the LLT node ID as 0 if LLT is not configured.

According to the protocol between the CP server and the application cluster, when you run the `cpsadm` on an application cluster node, `cpsadm` needs to send the LLT node ID of the local node to the CP server. But if LLT is unconfigured temporarily, or if the node is a single-node VCS configuration where LLT is not configured, then

the cpsadm command cannot retrieve the LLT node ID. In such situations, the cpsadm command fails.

Workaround: Set the value of the CPS_NODEID environment variable to 255. The cpsadm command reads the CPS_NODEID variable and proceeds if the command is unable to get LLT node ID from LLT.

Server-based fencing comes up incorrectly if default port is not mentioned (2403453)

When you configure fencing in customized mode and do no provide default port, fencing comes up. However, the vxfenconfig -l command output does not list the port numbers.

Workaround: Retain the "port=<port_value>" setting in the /etc/vxfenmode file, when using customized fencing with at least one CP server. The default port value is 14250.

Unable to customize the 30-second duration (2551621)

When the vxcpserv process is not able to bind to an IP address during startup, it attempts to bind to that IP address at an interval of 30 seconds. This interval is not configurable.

Workaround: There is no workaround for this issue.

NIC resource gets created with incorrect name while configuring CPSSG with the configure_cps.pl script (2585229)

The name of the NIC resource created by the configure_cps.pl script does not come out correct when, for example, m^{th} VIP is mapped to n^{th} NIC and every m is not equal to n. In this case, although CPSSG continues to function without any problem, when you unconfigure CPSSG using configure_cps.pl, it fails.

Workaround: To unconfigure CPSSG, you must remove the CPSSG configuration from the VCS configuration.

The cpsadm command fails after upgrading CP server to 6.0 or above in secure mode (2846727)

The cpsadm command may fail after you upgrade coordination point server (CP server) to 6.0 in secure mode. If the old VRTSat RPM is not removed from the system, the cpsadm command loads the old security libraries present on the system. As the installer runs the cpsadm command on the CP server to add or upgrade the SVS cluster (application cluster), the installer also fails.

Workaround: Perform the following procedure on all of the nodes of the CP server.

To resolve this issue

1 Rename `cpsadm` to `cpsadmbin`:

```
# mv /opt/VRTScps/bin/cpsadm /opt/VRTScps/bin/cpsadmbin
```

2 Create a file `/opt/VRTScps/bin/cpsadm` with the following content:

```
#!/bin/sh
EAT_USE_LIBPATH="/opt/VRTScps/lib"
export EAT_USE_LIBPATH
/opt/VRTScps/bin/cpsadmbin "$@"
```

3 Change the permissions of the new file to 775:

```
# chmod 755 /opt/VRTScps/bin/cpsadm
```

CoordPoint agent does not report the addition of new disks to a Coordinator disk group [2727672]

The LevelTwo monitoring of the CoordPoint agent does not report a fault even if the constituent of a coordinator disk group changes due to addition of new disks in the cooridnator disk group

Workaround: There is no workaround for this issue.

Coordination point server-based fencing may fail if it is configured on 5.1SP1RP1 using 6.0.1 coordination point servers (2824472)

The 5.1SP1 installer (CPI) cannot set up trust between a 5.1SP1 client and a 6.0 or later server, because there are no separate directories for truststores in the 5.1SP1. When trust cannot be setup, the 5.1SP1 installer cannot configure 5.1SP1 clients to work with 6.0 or later CPS in secure mode.

Workaround:

Set up trust manually between the CPS and clients using the cpsat or the vcsat command. After that, CPS and client will be able to communicate properly in the secure mode.

The upper bound value of FaultTolerance attribute of CoordPoint agent should be less than the majority of the coordination points. (2846389)

The upper bound value of `FaultTolerance` attribute of `CoordPoint` agent should be less than the majority of the coordination points. Currently this value is less than the number of coordination points.

Hostname and username are case sensitive in CP server (2846392)

The hostname and username on the CP server are case sensitive. The hostname and username used by fencing to communicate with CP server must be in same case as present in CP server database, else fencing fails to start.

Workaround: Make sure that the same case is used in the hostname and username on the CP server.

Virtual machine may return the not-responding state when the storage domain is inactive and the data center is down (2848003)

In a Red Hat Enterprise Virtualization Environment, if the storage domain is in an inactive state and the data center is in down state, the virtual machine may return a not-responding state and the KVMGuest resource in OFFLINE state.

Workaround: To resolve this issue:

1. Activate the storage domain in RHEV-M.
2. Check that the data center is in the up state.

Cannot run the vxfentsthdw utility directly from the install media if VRTSvxfen package is not installed on the system (2858190)

If VRTSvxfen package is not installed on the system, then certain script files that are needed for the vxfentsthdw utility to function are not available. So, without the VRTSvxfen package installed on the system you cannot run the utility from the install media.

Workaround: Install VRTSvxfen package, then run the utility from either the install media or from the `/opt/VRTSvcs/vxfen/bin/` location.

Server-based fencing may fail to start after reinstalling the stack (2802682)

Server-based fencing may fail to start if you use the existing configuration files after reinstalling the stack.

Workaround:

After reinstalling the stack, add the client cluster information on the coordination point server because the client cluster information is removed when the stack is uninstalled. For more details, see the Setting up server-based I/O Fencing manually section in the Symantec VirtualStore Installation Guide. Alternatively, you can manually modify the `/etc/vxfenmode` file and the `main.cf` file to start fencing in disable mode and then configure fencing.

Installation known issues

This section describes the known issues during installation and upgrade.

EULA changes (2161557)

The locations for all EULAs have changed.

The English EULAs now appear in /*product_dir*/EULA/en/EULA_*product_platform_version*.pdf

The EULAs for Japanese and Chinese now appear in those language in the following locations:

The Japanese EULAs appear in /*product_dir*/EULA/ja/EULA_*product_platform_version*.pdf

The Chinese EULAs appear in /*product_dir*/EULA/zh/EULA_*product_platform_version*.pdf

The VRTSacclib RPM is deprecated (2032052)

The VRTSacclib RPM is deprecated. For installation, uninstallation, and upgrades, note the following:

- Fresh installs: Do not install VRTSacclib.
- Upgrade: Ignore VRTSacclib.
- Uninstall: Ignore VRTSacclib.

Erroneous resstatechange trigger warning

You may encounter the following warning when you restart resources:

```
CPI WARNING V-9-40-4317 The installer has detected that resstatechange
trigger is configured by setting TriggerResStateChange attributes.
```

Workaround: In future releases, the resstatechange trigger will not be invoked when a resource is restarted. Instead, the resrestart trigger will be invoked if you set the TriggerResRestart attribute. The resrestart trigger is available in the current release. Refer to the VCS documentation for details.

Node is not able to join the cluster in case of full storage failure if _volasym tunable is on and mixed mode/disk based fencing is configured.(2755786)

Workaround:

The Web installer hangs at the end of the rolling upgrade process (2792835)

At the end of a rolling upgrade, the Web installer completes all the processes successfully but does not show the completion page.

Workaround:

Even though you don't see a completion page, the upgrade process executes successfully. Refresh the browser to begin using it for other purposes.

The uninstaller does not remove all scripts (2696033)

After removing SVS, some of the RC scripts remain in the `/etc/rc*.d/` folder. This is due to an issue with the chkconfig rpm in RHEL6 and updates. You can manually remove the scripts from the `/etc/rc*.d/` folder after removing the VxVM packages.

Workaround:

Install the chkconfig-1.3.49.3-1 chkconfig rpm from the RedHat portal. Refer to the following links:

http://grokbase.com/t/centos/centos/117pfhe4zz/centos-6-0-chkconfig-strange-behavior

http://rhn.redhat.com/errata/RHBA-2012-0415.html

Stopping the Web installer causes Device Busy error messages (2633924)

If you start the Web installer, and then perform an operation (such as prechecking, configuring, or uninstalling), you may get an error message saying the device is busy.

Workaround: Do one of the following:

- Kill the start.pl process.
- Start the webinstaller again. On the first Web page you see that the session is still active. Either take over this session and finish it or terminate it directly.

Ignore certain errors after an operating system upgrade—after a product upgrade with encapsulated boot disks (2030970)

Ignore certain errors after an operating system upgrade after a product upgrade with encapsulated boot disks.

You can ignore the following errors after you upgrade the operating system after a product upgrade that occurred with an encapsulated boot disk. Examples of the errors follow:

```
The partioning on disk /dev/sda is not readable by
The partioning tool parted, which is used to change the
partition table.
You can use the partitions on disk /dev/sda as they are.
You can format them and assign mount points to them, but you
cannot add, edit, resize, or remove partitions from that
disk with this tool.
```

Or

```
Root device: /dev/vx/dsk/bootdg/rootvol (mounted on / as reiserfs)
Module list: pilix mptspi qla2xxx silmage processor thermal fan
reiserfs aedd (xennet xenblk)

Kernel image; /boot/vmlinuz-2.6.16.60-0.54.5-smp
Initrd image: /boot/initrd-2.6.16.60-0.54.5-smp
```

The operating system upgrade is not failing. The error messages are harmless.

Workaround: Remove the /boot/vmlinuz.b4vxvm and /boot/initrd.b4vxvm files (from an un-encapsulated system) before the operating system upgrade.

After finishing a kernel upgrade on a master node the cvm group on a slave node does not come online (2439439)

After successfully finishing a kernel upgrade on one node, the cvm group does not come online on the second node.

Workaround: Check that your cluster is not in a jeopardy state before you perform a rolling upgrade.

Web installer does not ask for authentication after the first session if the browser is still open (2509330)

If you install or configure SVS and then close the Web installer, if you have other browser windows open, the Web installer does not ask for authentication in the subsequent sessions. Since there is no option to log out of the Web installer, the session remains open as long as the browser is open on the system.

Workaround: Make sure that all browser windows are closed to end the browser session and subsequently log in again.

Upgrading from Veritas Storage Foundation 5.1 Service Pack 1 Rolling Patch 2 to 6.0.1 with rootability enabled fails (2581313)

Upgrading from Veritas Storage Foundation (SF) 5.1 Service Pack (SP) 1 Rolling Patch (RP) 2 to 6.0.1 while using an encapsulated root disk fails because the post installation scripts of Veritas Volume Manager (VxVM) are unable to start the `initrd` daemon.

Workaround: To upgrade from 5.1 SP1 RP2 to 6.0.1 while using an encapsulated root disk, you must reinstall the nash utility on the system prior to the upgrade.

To upgrade from 5.1 SP1 RP2 to 6.0.1 while using an encapsulated root disk

1. Encapsulate the root disk.
2. Reinstall the nash utility.
3. Upgrade to the SF 6.0.1 release.

Stopping the installer during an upgrade and then resuming the upgrade might freeze the service groups [2574731]

The service groups freeze due to upgrading using the product installer if you stopped the installer after the installer already stopped some of the processes and then resumed the upgrade.

Workaround:

You must unfreeze the service groups manually after the upgrade completes.

To unfreeze the service groups manually

1 List all the frozen service groups

```
# hagrp -list Frozen=1
```

2 Unfreeze all the frozen service groups:

```
# haconf -makerw
# hagrp -unfreeze service_group -persistent
# haconf -dump -makero
```

After a locale change restart the vxconfig daemon (2417547)

You need to restart the vxconfig daemon you change the locale of nodes that use it. The vxconfig daemon starts at boot. If you have changed locale, you need to restart the daemon.

Workaround: See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more information on vxconfigd daemon recovery.

Error messages in syslog (1630188)

If you install or uninstall a product on a node, you may see the following warnings in syslog: /var/log/message. These warnings are harmless and can be ignored.

```
Jul  6 10:58:50 swlx62 setroubleshoot: SELinux is preventing the
semanage from using potentially mislabeled files
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For
complete SELinux messages. run sealert -l ed8978d1-0b1b-4c5b-a086-
67da2a651fb3
Jul  6 10:58:54 swlx62 setroubleshoot: SELinux is preventing the
semanage from using potentially mislabeled files
(/var/tmp/installer-200907061052eVe/install.swlx62.VRTSvxvm). For
complete SELinux messages. run sealert -l ed8978d1-0b1b-4c5b-a086-
67da2a651fb3
Jul  6 10:58:59 swlx62 setroubleshoot: SELinux is preventing the
restorecon from using potentially mislabeled files
```

NetBackup 6.5 or older version is installed on a VxFS file system (2056282)

If you have NetBackup 6.5 or older version installed on a VxFS file system and before upgrading to Veritas Storage Foundation (SF) 6.0.1, if you unmount all VxFS

file systems including the one that hosts the NetBackup binaries (/usr/openv), then while upgrading to SF 6.0.1, the installer fails to check if NetBackup is installed on the same machine and uninstalls the shared infrastructure RPMs VRTSpbx, VRTSat, and VRTSicsco. This causes NetBackup to stop working.

Workaround: Before you unmount the VxFS file system that hosts NetBackup, copy the /usr/openv/netbackup/bin/version file and /usr/openv/netbackup/version file to the /tmp directory. If you have clustered NetBackup installed, you must also copy the /usr/openv/netbackup/bin/cluster/NBU_RSP file to the /tmp directory. After you unmount the NetBackup file system, manually copy these two version files from /tmp to their original directories. If you have clustered NetBackup installed, you must also copy the /usr/openv/netbackup/bin/cluster/NBU_RSP file from /tmp to its original directory.

If the version files' directories do not exist, create the directories:

```
# mkdir -p /usr/openv/netbackup/bin
```

Run the installer to finish the upgrade process. After upgrade process completes, remove the two version files and their directories.

If your system is already affected by this issue, then you must manually install the VRTSpbx, VRTSat, and VRTSicsco RPMs after the upgrade process completes.

During upgrade from 5.1SP1 with an encapsulated root disk, splitting the root mirror fails if the target disk group name is used by a deported disk group (2280560)

During an upgrade from SVS 5.1 SP1 with an encapsulated root disk, splitting the root mirror fails if the target disk group name for the split operation is used by an existing deported disk group.

Workaround:

Specify a different disk group name as a target for the split operation.

Installer is unable to split a cluster that is registered with one or more CP servers (2110148)

Splitting a cluster that uses server-based fencing is currently not supported.

You can split a cluster into two and reconfigure Symantec VirtualStore HA on the two clusters using the installer. For example, you can split a cluster *clus1* into *clus1A* and *clus1B*.

However, if you use the installer to reconfigure the Symantec VirtualStore HA, the installer retains the same cluster UUID of *clus1* in both *clus1A* and *clus1B*. If both *clus1A* and *clus1B* use the same CP servers for I/O fencing, then the CP server allows registration only from the cluster that attempts to register first. It rejects the registration from the cluster that attempts next. Thus, the installer reports failure during the reconfiguration of the cluster that uses server-based fencing.

Workaround: There is no workaround for this issue.

Veritas File System modules fail to unload during uninstall or upgrade if a break-off snapshot volume is created or reattached (2851403)

If a break-off snapshot volume is created or reattached on the system, the Veritas File System modules, `vxportal` and `vxfs`, may fail to unload during uninstall or upgrade. The situation occurs if the SmartMove feature is enabled, which is the default setting. When you use the installer to uninstall or upgrade, you may see a message similar to the following:

```
Veritas Storage Foundation Shutdown did not complete successfully

vxportal failed to stop on dblxx64-21-v1
vxfs failed to stop on dblxx64-21-v1
```

Workaround:

1. Open a new session and manually unload the modules that failed to unload. Use commands similar to the following:

   ```
   # /sbin/modprobe -r vxportal
   # /sbin/modprobe -r vxfs
   ```

2. Because some processes failed to stop, the installer recommends a reboot and asks you if you want to continue.

   Press `y` to continue to the next phase. You can ignore the reboot requirement.

Software limitations

The following are software limitations in the 6.0.4 release of Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

Thin reclamation requests are not redirected when the ioship policy is enabled

Reclamation requests fail from nodes that do not have connectivity to the disks, even when the ioship policy is enabled. Reclamation I/Os are not redirected to another node.

VMware vSphere extension for VirtualStore limitations

The following are the software limitations for VMware vSphere extension for VirtualStore that are known in this release.

F5 usage is not supported for wizard refreshing (2362940)

F5 usage is not supported for wizard refreshing.

Workaround

To get new or refreshed data, it is important to restart the wizard and not use the F5 key.

Virtual machines with VMware Snapshots cannot be used as golden images (2514969)

Any virtual machine (or template) which has VMware Snapshots stored, cannot be used as a golden image for making clones with the FileSnap wizard. To use such virtual machines (or templates), first delete the Snapshots, then use the FileSnap wizard.

Replication software limitations

The following are replication software limitations in this release of Symantec VirtualStore.

Softlink access and modification times are not replicated on SLES10 for VFR jobs

When running a file replication job on SLES10, softlink access and modification times are not replicated.

Limitations related to I/O fencing

This section covers I/O fencing-related software limitations.

Preferred fencing limitation when VxFEN activates RACER node re-election

The preferred fencing feature gives preference to more weighted or larger subclusters by delaying the smaller subcluster. This smaller subcluster delay is effective only if the initial RACER node in the larger subcluster is able to complete the race. If due to some reason the initial RACER node is not able to complete the race and the VxFEN driver activates the racer re-election algorithm, then the smaller subcluster delay is offset by the time taken for the racer re-election and the less weighted or smaller subcluster could win the race. This limitation though not desirable can be tolerated.

Uninstalling VRTSvxvm causes issues when VxFEN is configured in SCSI3 mode with dmp disk policy (2522069)

When VxFEN is configured in SCSI3 mode with dmp disk policy, the DMP nodes for the coordinator disks can be accessed during system shutdown or fencing arbitration. After uninstalling VRTSvxvm RPM, the DMP module will no longer be loaded in memory. On a system where VRTSvxvm RPM is uninstalled, if VxFEN attempts to access DMP devices during shutdown or fencing arbitration, the system panics.

Documentation

Product guides are available in the PDF format on the software media in the `/docs/product_name` directory. Additional documentation is available online.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The latest product documentation is available on the Symantec website.

http://sort.symantec.com/documents

Documentation set

Table 1-8 lists the documentation for Veritas Storage Foundation Cluster File System High Availability.

Table 1-8 Veritas Storage Foundation Cluster File System High Availability documentation

Document title	File name
Veritas Storage Foundation Cluster File System High Availability Release Notes	sfcfs_notes_604_lin.pdf
Veritas Storage Foundation Cluster File System High Availability Installation Guide	sfcfs_install_604_lin.pdf
Veritas Storage Foundation Cluster File System High Availability Administrator's Guide	sfcfs_admin_604_lin.pdf

Table 1-9 lists the documentation for Symantec VirtualStore.

Table 1-9 Symantec VirtualStore documentation

Document title	File name
Symantec VirtualStore Release Notes	virtualstore_notes_604_lin.pdf
Symantec VirtualStore Installation and Configuration Guide	virtualstore_install_604_lin.pdf
Symantec VirtualStore Administrator's Guide	virtualstore_admin_604_lin.pdf

If you use Veritas Operations Manager (VOM) to manage Veritas Storage Foundation and High Availability products, refer to the VOM product documentation at:

http://sort.symantec.com/documents

Manual pages

The manual pages for Veritas Storage Foundation and High Availability Solutions products are installed in the `/opt/VRTS/man` directory.

Set the `MANPATH` environment variable so the `man`(1) command can point to the Veritas Storage Foundation manual pages:

- For the Bourne or Korn shell (`sh or ksh`), enter the following commands:

```
MANPATH=$MANPATH:/opt/VRTS/man
  export MANPATH
```

- For C shell (`csh` or `tcsh`), enter the following command:

```
setenv MANPATH ${MANPATH}:/opt/VRTS/man
```

See the man(1) manual page.

Manual pages are divided into sections 1, 1M, 3N, 4, and 4M. Edit the man(1) configuration file /etc/man.config to view these pages.

To edit the man(1) configuration file

1. If you use the man command to access manual pages, set LC_ALL to “C” in your shell to ensure that the pages are displayed correctly.

   ```
   export LC_ALL=C
   ```

   See incident 82099 on the Red Hat Linux support website for more information.

2. Add the following line to /etc/man.config:

   ```
   MANPATH /opt/VRTS/man
   ```

   where other man paths are specified in the configuration file.

3. Add new section numbers. Change the line:

   ```
   MANSECT         1:8:2:3:4:5:6:7:9:tcl:n:l:p:o
   ```

   to

   ```
   MANSECT         1:8:2:3:4:5:6:7:9:tcl:n:l:p:o:3n:1m
   ```

The latest manual pages are available online in HTML format on the Symantec website at:

https://sort.symantec.com/documents