

Veritas Storage Foundation™ and High Availability Solutions 5.0 Maintenance Pack 2 Rolling Patch 1

Introduction

This document provides release information about Rolling Patch 1 for the Veritas Storage Foundation and High Availability Solutions 5.0 Maintenance Pack 2 release on the Linux platform. The Rolling Patch adds kernel support for SUSE Linux Enterprise Server 9 Service Pack 4. For Veritas Storage Foundation, this Rolling Patch also includes fixed issues.

Supported Platforms

Rolling Patch 1 (RP1) for Veritas Storage Foundation 5.0 Maintenance Pack 2 can only be installed on the following Linux operating systems and kernel binaries distributed by Oracle, Red Hat and SUSE:

Operating System	Architecture	Kernel
Red Hat Enterprise Linux 4 Update 3	x86_64 (Intel Xeon, AMD Opteron)	2.6.9-34 kernels
Red Hat Enterprise Linux 4 Update 4	x86_64 (Intel Xeon, AMD Opteron)	2.6.9-42 kernels
SUSE Linux Enterprise Server 9 Service Pack 3	x86_64 (Intel Xeon, AMD Opteron)	2.6.5-7.244, 252, 276 kernels
SUSE Linux Enterprise Server 9 Service Pack 4	x86_64 (Intel Xeon, AMD Opteron)	2.6.5-7.308 kernels
Oracle Enterprise Linux (based on, and has binary compatibility with, Red Hat Enterprise Linux 4 Update 4)	x86_64 (Intel Xeon, AMD Opteron)	2.6.9-42 kernels

Rolling Patch 1 (RP1) for Veritas Storage Foundation and High Availability Solutions 5.0 Maintenance Pack 2 can only be installed on the following Linux operating systems and kernel binaries distributed by SUSE:

Operating System	Architecture	Kernel
SUSE Linux Enterprise Server 9 Service Pack 4	x86 (32-bit)	2.6.5-7.308-smp kernel 2.6.5-7.308-bigsm kernel 2.6.5-7.308-default kernel
	x86_64 (Intel Xeon, AMD Opteron)	2.6.5-7.308-default kernel 2.6.5-7.308-smp kernel

Veritas products will operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

Information about the latest supported Red Hat erratas and updates and SUSE service packs is available in the following TechNote. Read this TechNote *before* installing any Veritas™ product.

<http://support.veritas.com/docs/281993>

For further details, see the *Veritas Storage Foundation 5.0 MP2 Release Notes*. and the *Veritas Cluster Server 5.0 MP2 Release Notes*.

Warning Storage Foundation supports SUSE Linux Enterprise Server 9 Service Pack 4 with this Rolling Patch release. However, customers running EMC PowerPath should not upgrade to SUSE Linux Enterprise Server 9 Service Pack 4 until it is officially supported by EMC.

Fixed Issues

The following Veritas Volume Manager (VxVM) incidents are fixed in this Rolling Patch:

Incident	Description
999019	Errors related to duplicate disk identifier seen while importing a disk group with underlying disks that are hardware mirrored.
1003063	Panic in voliod_iohandle() as it tries to use previously freed sio.
1003067	When enabling the third disabled path, DMP nodes with array EVA8000 get disabled.
1004009	vxdisksetup sizes EFI LUNs incorrectly.
1004051	Growing a volume with a DCO association causes vxconfigd to hang during transaction commit.
1004131	Volume Manager nodes in a cluster remain in 'master selection' after random rebooting with the 'shutdown -r now' command.
1004319	The vxcvmdiag cvminfo command aborts with an error.
1009712	Node join fails with the 'recovery in progress' error in a Clustered Volume Manager dirty region logging (DRL) configuration.
1013211	LUN reuse issue that was not fixed by the DMP backport hotfix
1013220	vxdisk updateudid after udid_mismatch leads to disk group configuration issues.
1036602	dmp_get_path_state() has been modified to check only for write protection on Symmetrix arrays.
1036605	PID in CLARiiON ASL has been corrected.
1040229	vxvoltune cannot be used to set maxiosize to a value greater than 512.
1046715	VxVM commands take unexpected amount of time with CLARiiON.
1046735	Restores fail for raw partition tablespaces.
1053489	Read data corruption occurs during I/O to a volume.
1062775	Panic occurs in the dmplinux_close_disk_device() function.
1065339	Memory space errors because the memory used by the Volume Manager and dynamic multipathing (DMP) plug-ins of VERITAS Federated Mapping Service (VxMS) is not freed.



Incident	Description
1066677	The vxdmp device driver for VxMS does not return all the paths on Linux.
1066784	Errors seen when the Volume Manager plugin of VxMS attempts to use object names with storage extensions.
1066793	Volume Manager plugin of VxMS fails to initialize the private data of log plex when a volume is having DRL log enabled, causing buildmapdomain to fail.
1066798	Added support for passive paths.
1066804	dg_set_current fails for bootdg (an alias for the boot disk group).
1070085	NetApp takeover and giveback result in disk or disk group failures.
1070977	Backup fails as VxVM assigns the same DMP node to two NetApp LUNs.
1093584	The 'vxassist grow' command run on a volume with a fragmented subdisks layout and more than two log plexes per sub-volume results in an error, though the 'vxassist maxgrow' command reports that the volume can be grown.
1093586	When converting a vxprint dumpconfig to a vxprint makefile, it fails because a variable contains an unusable value.
1093593	Changes made to the vxrecover command to handle multiple recoveries.
1093602	Import of a disk group fails due to a flag being incorrectly set in da_dg_import_fixup.
1093621	System panics with a data page fault at vxtask_created() while mirroring a VxVM volume.
1093624	System panics when a "differences based sync" is performed on a replicated volume group.
1093628	System hangs on a primary host with heavy I/O load.
1093629	Added support for mixed naming schemes.
1093631	Oracle loops endlessly due to broken VOL_READ_MIRRORS ioctl behavior.
1093635	False serial split brain detection.
1094091	During master failover, after rebooting the master, all nodes go in kernel debugger (kdb).
1094277	Full resync occurs in the remaining nodes after SFORAC panic reboots.
1094287	Data corruption due to DCO maps that get cleared unexpectedly.
1094367	The vxassist relay command followed by the vxconvert command disturbs the HA configuration leading to the plexes of a mirror being disabled on controller failure.
1094375	System hangs during replication.
1095967	Failover mode support added for CLARiiON arrays.
1096038	VxVM stack-handoff code works incorrectly.
1098940	Panic occurs in voldco_update_instant.
1103671	Erroneous message seen due to a spurious check in the post-install script.
1105739	Enabled initialization of LUNs that are greater than 2TB.
1172421	Node in a cluster panics when the array is disconnected.
1173291	CVM master node panics when the HBA is disconnected.



Incident	Description
1173302	CVM master node panics in the <code>spinlocks_held_leaving_processor()</code> function.

The following Veritas File System (VxFS) incidents are fixed in this Rolling Patch:

Incident	Description
990192	<code>fsmap -N</code> is accepted on Linux, but does not show named streams.
990194	Named data stream support on Linux is inadequate.
1018537	<code>fsck</code> calls <code>vol_dereserve()</code> with the wrong argument.
1029541	<code>VX_GET_NINODE</code> value changed back to 57.
1037211	Panic during db2 fast pre-allocation with VxFS.
1042774	<code>vx_write</code> performance is impacted when full FS.
1052736	VxFS's ACL invalid when accessed via Samba or Windows (group information disappears).
1060159	Mapping provider returns incorrect directory entries.
1069075	Several <code>vx_dotdot_manipulate</code> messages appear after resize and <code>vxupgrade</code> .
1084159	Non-root user cannot run <code>vxquota</code> to get the group quota setting.
1085224	<code>vx_getattr()</code> has been enhanced to take <code>IGLOCK</code> in share mode.
1093424	<code>vx_logbuf_clean</code> sleeps while holding lock leading to performance degradation.
1094392	FS hangs when it tries to free a large extent.
1097091	Poor read performance through FS if the file is in use by ODM or QIO.
1097094	Data page fault in <code>vx_cbuf_revoke_blk</code> .
1098186	<code>fsck</code> log replay was slow.
1100079	<code>vxdump</code> incorrectly estimates size of tape device.
1100090	Redundant <code>VX_DELAY</code> in <code>vx_glock_grant</code> and <code>vx_rwlock_range_grant</code> that leads to CFS performance degradation.
1100091	<code>vxupgrade</code> from Layout5 to Layout6 was resulting in inode table overflow.
1100092	<code>vx_ireuse</code> han stall causes CFS hang.

Downloading the RP1 Archive

The patches comprising the RP1 release are available for download from the Veritas website. After downloading the RP1 file, use the `tar` command to uncompress and extract the archive.



Packages Included in this Rolling Patch

The following packages are included in this rolling patch for Storage Foundation on Red Hat Enterprise Linux:

Operating System	Arch.	Packages
Red Hat Enterprise Linux 4	x86_64	VRTSvxvm-common-5.0.20.10-MP2RP1_RHEL4.i686.rpm VRTSvxvm-platform-5.0.20.10-MP2RP1_RHEL4.x86_64.rpm VRTSvxfs-common-5.0.20.10-MP2RP1_RHEL4.i686.rpm VRTSvxfs-platform-5.0.20.10-MP2RP1_RHEL4.x86_64.rpm

The following packages are included in this rolling patch for Storage Foundation on SUSE Linux Enterprise Server:

Operating System	Arch.	Packages
SUSE Linux Enterprise Server 9	x86_64	VRTSvxvm-common-5.0.20.10-MP2RP1_SLES9.i586.rpm VRTSvxvm-platform-5.0.20.10-MP2RP1_SLES9.x86_64.rpm VRTSvxfs-common-5.0.20.10-MP2RP1_SLES9.i586.rpm VRTSvxfs-platform-5.0.20.10-MP2RP1_SLES9.x86_64.rpm VRTSllt-5.0.20.10-MP2RP1_SLES9.x86_64.rpm VRTSgab-5.0.20.10-MP2RP1_SLES9.x86_64.rpm VRTSvxfen-5.0.20.10-MP2RP1_SLES9.x86_64.rpm VRTSvcsdr-5.0.20.10-MP2RP1_SLES9.x86_64.rpm VRTSg1m-5.0.20.10-MP2RP1a_SLES9.x86_64.rpm VRTSgms-5.0.20.10-MP2RP1a_SLES9.x86_64.rpm VRTSodm-common-5.0.20.10-MP2RP1_SLES9.i586.rpm VRTSodm-platform-5.0.20.10-MP2RP1_SLES9.x86_64.rpm

The following packages are included in this rolling patch for Storage Foundation and High Availability Solutions on SUSE Linux Enterprise Server:

Operating System	Arch.	Packages
SUSE Linux Enterprise Server 9	x86 (32-bit)	VRTSgab-5.0.20.10-MP2RP1_SLES9.i586.rpm VRTSvcsdr-5.0.20.10-MP2RP1_SLES9.i586.rpm VRTSllt-5.0.20.10-MP2RP1_SLES9.i586.rpm VRTSvxfen-5.0.20.10-MP2RP1_SLES9.i586.rpm
SUSE Linux Enterprise Server 9	x86_64	VRTSgab-5.0.20.10-MP2RP1_SLES9.x86_64.rpm VRTSvcsdr-5.0.20.10-MP2RP1_SLES9.x86_64.rpm VRTSllt-5.0.20.10-MP2RP1_SLES9.x86_64.rpm VRTSvxfen-5.0.20.10-MP2RP1_SLES9.x86_64.rpm

The following sections contain instructions for installing this rolling patch on a cluster and on a standalone system.



Installing RP1 on a Cluster

An upgrade requires stopping cluster failover functionality during the entire procedure, however CFS and CVM services remain available. The upgrade is performed in several stages:

- ◆ Freeze service group operations and stop VCS on the cluster.
- ◆ Select a group of one or more cluster nodes to upgrade, and leave a group of one or more nodes running.
- ◆ Take the first group offline and install the software patches.
- ◆ Bring the first group (with the newly installed patches) online to restart cluster failover services.
- ◆ Upgrade the nodes in the second group, and bring them online. The cluster is fully restored.

To install RP1 on a cluster:

1. Log in as superuser.
2. Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
3. Switch the service group to a node that is running.

```
# hagrps -switch service_group -to nodename
```
4. From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```
5. Enter the following command to freeze HA service group operations on each node:

```
# hasys -freeze -persistent nodename
```
6. Make the configuration read-only:

```
# haconf -dump -makero
```
7. Select the group of nodes that are to be upgraded first, and follow [step 8](#) through [step 33](#) for these nodes.
8. Stop VCS by entering the following command on each node in the group being upgraded:

```
# hastop -local
```
9. Stop the VCS command server:

```
# killall CmdServer
```
10. Unregister CFS from GAB, and stop cluster fencing, GAB and LLT:

```
# fsclustadm cfsdeinit  
# /etc/init.d/vxfen stop  
# /etc/init.d/gab stop  
# /etc/init.d/llt stop
```



-
11. Check if each node's root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- a. Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Note Do not remove the plexes on the root disk that correspond to the original disk partitions.

- b. Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

12. If required, you can upgrade the nodes at this stage, and patch them to a supported kernel version.

13. On each node, use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

14. On each node in the cluster, unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

15. On each node, verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | fsdb -t vxfs filesystem | grep clean  
flags 0 mod 0 clean clean_value
```

A `clean_value` value of `0x5a` indicates the file system is clean, `0x3c` indicates the file system is dirty, and `0x69` indicates the file system is dusty. A dusty file system has pending extended operations.

- a. If a file system is not clean, enter the following commands for that file system:

```
# fsck -t vxfs filesystem  
# mount -t vxfs filesystem mountpoint  
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.

There may be a pending large fileset clone removal extended operation if the `umount` command fails with the following error:

```
file system device busy
```



You know for certain that an extended operation is pending if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system
file system still in progress.
```

- b. If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large filesset clone can take several hours.
- c. Repeat [step 15](#) to verify that the unclean file system is now clean.

16. If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- a. Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- b. Use the `vxrvvg stop` command to stop each RVG individually:

```
# vxrvvg -g diskgroup stop rvg_name
```
- c. On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

Caution To avoid data corruption, do not proceed until all RLINKs are up-to-date.

17. Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

18. On each node, stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

To verify that no volumes remain open, use the following command:

```
# vxprint -Aht -e v_open
```

19. Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

20. On each node, use the following commands to upgrade to 5.0 MP2 RP1.

```
# killall vxpal
# rpm -U VRTSvxvm-common-5.0.20.10-MP2RP1_system.arch.rpm
# rpm -U VRTSvxvm-platform-5.0.20.10-MP2RP1_system.arch.rpm
# rpm -U VRTSvxfs-common-5.0.20.10-MP2RP1_system.arch.rpm
# rpm -U VRTSvxfs-platform-5.0.20.10-MP2RP1_system.arch.rpm
```

where *system* is RHEL4 or SLES9, and *arch* is i586, i686 or x86_64 as appropriate. (See [“Packages Included in this Rolling Patch”](#) on page 5.)



If you are upgrading a SUSE Linux Enterprise Server 9 Service Pack 4 node that has Veritas Storage Foundation Cluster File System (SFCFS), or SFCFS-RAC, installed on it, additionally run the following commands:

```
# rpm -U VRTSllt-5.0.20.10-MP2RP1_SLES9.x86_64.rpm
# rpm -U VRTSgab-5.0.20.10-MP2RP1_SLES9.x86_64.rpm
# rpm -U VRTSvxfen-5.0.20.10-MP2RP1_SLES9.x86_64.rpm
# rpm -U VRTSvcldr-5.0.20.10-MP2RP1_SLES9.x86_64.rpm
# rpm -U VRTSglm-5.0.20.10-MP2RP1a_SLES9.x86_64.rpm
```

On a SUSE Linux Enterprise Server 9 Service Pack 4 node that has SFCFS-RAC installed on it, additionally run the following commands:

```
# rpm -U VRTSgms-5.0.20.10-MP2RP1a_SLES9.x86_64.rpm
# rpm -U VRTSodm-common-5.0.20.10-MP2RP1_SLES9.i586.rpm
# rpm -U VRTSodm-platform-5.0.20.10-MP2RP1_SLES9.x86_64.rpm
```

If you are upgrading a SUSE Linux Enterprise Server 9 Service Pack 4 node that has Veritas Cluster Server installed on it, additionally run the following commands:

```
# rpm -U VRTSllt-5.0.20.10-MP2RP1_SLES9.<arch>.rpm
# rpm -U VRTSgab-5.0.20.10-MP2RP1_SLES9.<arch>.rpm
# rpm -U VRTSvxfen-5.0.20.10-MP2RP1_SLES9.<arch>.rpm
# rpm -U VRTSvcldr-5.0.20.10-MP2RP1_SLES9.<arch>.rpm
```

21. Shut down and reboot each of the upgraded nodes. After the nodes come back up, application failover capability is available for that group.
22. If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.
23. If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.
24. Start LLT, GAB, VXFEN, and VCS:

```
# /etc/init.d/llt start
# /etc/init.d/gab start
# /etc/init.d/vxfen start
# /etc/init.d/vcs start
```

25. Make the VCS configuration writable again from any node in the upgraded group:

```
# haconf -makerw
```

26. Enter the following command on each node in the upgraded group to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent nodename
```

27. Make the configuration read-only:

```
# haconf -dump -makero
```

28. Switch the service group to the original node:

```
# hagrps -switch service_group -to nodename
```

29. Bring the CVM service group online on each node in the upgraded group:



```
# hagrps -online cvm -sys nodename
```

30. Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

31. If you stopped any RVGs in [step 16](#), restart each RVG:

```
# vxrvrg -g diskgroup start rvg_name
```

32. Remount all VxFS file systems and Storage Checkpoints on all nodes:

```
# mount /filesystem  
# mount /checkpoint_name
```

33. Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctrl start
```

34. Repeat [step 8](#) through [step 33](#) for the second group of nodes.



Installing RP1 on a Standalone System

To install RP1 on a standalone system:

1. Log in as superuser.
2. Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
3. Check if the root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- a. Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Note Do not remove the plexes on the root disk that correspond to the original disk partitions.

- b. Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

4. If required, you can upgrade the system at this stage, and patch it to a supported kernel version.
5. Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

6. Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

7. Verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | fsdb -t vxfs filesystem | grep clean  
flags 0 mod 0 clean clean_value
```

A `clean_value` value of `0x5a` indicates the file system is clean, `0x3c` indicates the file system is dirty, and `0x69` indicates the file system is dusty. A dusty file system has pending extended operations.



-
- a. If a file system is not clean, enter the following commands for that file system:

```
# fsck -t vxfs filesystem
# mount -t vxfs filesystem mountpoint
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.

There may be a pending large fileset clone removal extended operation if the `umount` command fails with the following error:

```
file system device busy
```

You know for certain that an extended operation is pending if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system
file system still in progress.
```

- b. If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
 - c. Repeat [step 7](#) to verify that the unclean file system is now clean.
8. If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- a. Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- b. Use the `vxrvg stop` command to stop each RVG individually:

```
# vxrvg -g diskgroup stop rvg_name
```

- c. On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

Caution To avoid data corruption, do not proceed until all RLINKs are up-to-date.

9. Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

10. Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

To verify that no volumes remain open, use the following command:

```
# vxprint -Aht -e v_open
```

11. Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctl stop
```



12. Use the following commands to upgrade to 5.0 MP2 RP1.

```
# killall vxpal
# rpm -U VRTSvxvm-common-5.0.20.10-MP2RP1_system.arch.rpm
# rpm -U VRTSvxvm-platform-5.0.20.10-MP2RP1_system.arch.rpm
# rpm -U VRTSvxfs-common-5.0.20.10-MP2RP1_system.arch.rpm
# rpm -U VRTSvxfs-platform-5.0.20.10-MP2RP1_system.arch.rpm
```

where *system* is RHEL4 or SLES9, and *arch* is i586, i686 or x86_64 as appropriate. (See “[Packages Included in this Rolling Patch](#)” on page 5.)

13. Shut down and reboot the system.

14. If necessary, reinstate any missing mount points in the `/etc/fstab` file.

15. Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

16. If you stopped any RVGs in [step 8](#), restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

17. Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem
# mount /checkpoint_name
```

18. Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

19. If you need to re-encapsulate and mirror the root disk, follow the procedures in the “Administering Disks” chapter of the *Veritas Volume Manager Administrator’s Guide*.

Verifying Software Versions

To list the Veritas packages installed on your system, enter the following command:

```
# rpm -qa | egrep VRTS
```



Removing the RP1 packages

Roll back of the RP1 packages to the release 5.0 MP2 version of the packages is not supported. It is recommended that you follow the steps in this section to remove all the installed Veritas packages, and then perform a complete reinstallation of the release 5.0 MP2 software.

To uninstall the Veritas software:

1. Log in as superuser.
2. Verify that `/opt/VRTS/bin` is in your PATH so you can execute all product commands.
3. Stop VCS along with all the resources. Then, stop the remaining resources manually:

```
# /etc/init.d/vcs stop
```
4. Stop the VCS command server:

```
# killall CmdServer
```
5. Uninstall VCS:

```
# cd /opt/VRTS/install
# ./uninstallvcs
```
6. If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

```
# rm /etc/vxfenmode
```
7. Check if the root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- a. Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Note Do not remove the plexes on the root disk that correspond to the original disk partitions.

- b. Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the `rootdg` disk group in addition to the root disk for `vxunroot` to succeed.

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.



-
8. Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

9. Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
# umount /filesystem
```

10. If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- a. Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- b. Use the `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

- c. On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

Caution To avoid data corruption, do not proceed until all RLINKs are up-to-date.

11. Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

12. Stop all VxVM volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

To verify that no volumes remain open, use the following command:

```
# vxprint -Aht -e v_open
```

13. Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

14. To shut down and remove the installed Veritas packages, use the appropriate command in the `/opt/VRTS/install` directory. For example, to uninstall the Storage Foundation or Veritas Storage Foundation *for DB2* packages, use the following commands:

```
# cd /opt/VRTS/install
# ./uninstallsf [-usersh]
```

You can use this command to remove the packages from one or more systems. The `-usersh` option is required if you are using the remote shell (RSH) rather than the secure shell (SSH) to uninstall the software simultaneously on several systems.



Note Provided that the remote shell (RSH) or secure shell (SSH) has been configured correctly, this command can be run on a single node of the cluster to install the software on all the cluster nodes.

After uninstalling the Veritas software, reinstall the release 5.0 MP2 software as described in the *Veritas Storage Foundation 5.0 MP2 Release Notes*.

Getting Help

For technical assistance, visit

http://www.symantec.com/enterprise/support/assistance_care.jsp and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Trademarks

Copyright © 2008 VERITAS Software Corporation. All rights reserved. Symantec, the Symantec logo, Veritas, and Veritas Storage Foundation are trademarks or registered trademarks of Symantec Software Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.



Symantec Corporation
20330 Steven Creek Blvd.
Cupertino, CA 95014

