
Red Hat Enterprise Linux 6

Technical Notes

Technical Release Documentation



Copyright © 2010 Red Hat.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux® is the registered trademark of Linus Torvalds in the United States and other countries.

Java® is a registered trademark of Oracle and/or its affiliates.

XFS® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

All other trademarks are the property of their respective owners.

1801 Varsity Drive
Raleigh, NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701

Abstract

Red Hat Enterprise Linux 6.0 Technical notes provide details on various features shipped in Red Hat Enterprise Linux 6.0, as well as all known issues of this release.

Technical Notes

1. Installer	2
1.1. Known Issues	3
2. Deployment	5
2.1. Known Issues	8
3. Virtualization	10
3.1. Known Issues	11
4. Storage and Filesystems	14
4.1. Technology Previews	15
4.2. Known Issues	16
5. Networking	19
5.1. Technology Previews	19
5.2. Known Issues	19
6. Clustering	20
6.1. Technology Previews	20
6.2. Known Issues	20
7. Authentication	21
7.1. Technology Previews	21
7.2. Known Issues	22
8. Security	22
8.1. Technology Previews	22
9. Devices and Device Drivers	23
9.1. Technology Previews	23
9.2. Known Issues	23
10. Kernel	24
10.1. Technology Previews	25
10.2. Known Issues	25
11. Development and Tools	29
11.1. Technology Previews	29
11.2. Known Issues	29
12. Desktop	30
12.1. Known Issues	30
A. Package Manifest	32
B. Large Package Updates	32
B.1. kernel	32
B.1.1. RHSA-2011:0883 - Important: kernel security and bug fix update	32
B.1.2. RHSA-2011:0498 - Important: kernel security, bug fix and enhancement update	35
B.1.3. RHSA-2011:0421 - Important: kernel security and bug fix update	39
B.1.4. RHSA-2011:0283 - Moderate: kernel security, bug fix and enhancement update	44
B.1.5. RHSA-2011:0007 - Important: kernel security and bug fix update	48
B.1.6. RHSA-2010:0842: Important: kernel security and bug fix update	52
C. Revision History	58
Index	58

1. Installer

The Red Hat Enterprise Linux installer (also known as anaconda) assists in the installation of Red Hat Enterprise Linux 6.

Installation on systems with multipath and non-multipath storage devices

Installation of Red Hat Enterprise Linux 6 on a system with multipath and non-multipath storage devices the automatic partitioning layout in the installer may create volume groups containing a mix of multipath and non-multipath devices, thus defeating the purpose of multipath storage.

Users are advised to either select only multipath or only non-multipath devices on the disk selection screen that appears after selecting automatic partitioning. Alternatively, users can select custom partitioning.

1.1. Known Issues

- *The following issue applies to IBM Power Systems only.*

anaconda will not create a new PReP boot partition on the root disk when performing a new Red Hat Enterprise Linux 6 installation on a system that contains existing PReP Boot partitions that need to be preserved. Consequently, the Power SMS boot manager will be unable to boot the new Red Hat Enterprise Linux 6 installation. To work around this issue,

1. Use the `fdisk` utility to temporarily change the partition type from type 41 'PReP Boot' to type 83 'Linux' for all existing Linux installations on the system.
2. Perform the Red Hat Enterprise Linux 6 installation. During installation, a new PReP Boot partition will be created on the Red Hat Enterprise Linux 6 root disk.
3. Post-installation, once the new Red Hat Enterprise Linux 6 installation is up and running, use the `fdisk` utility to restore all changed partition types to type 41 'PReP Boot'.

- Anaconda now utilizes **NetworkManager** for network interface configuration. Consequently, kickstart users that referenced the network settings located in `/tmp/netinfo` must now source the **ifcfg** files found in `/etc/sysconfig/network-scripts`
- In some circumstances, disks that contain a whole disk format (e.g. a LVM Physical Volume populating a whole disk) are not cleared correctly using the `clearpart --initlabel` kickstart command. Adding the `--all` switch — as in `clearpart --initlabel --all` — ensures disks are cleared correctly.
- The `nodmraid` boot parameter currently cannot be used to force installation on disks containing spurious BIOS RAID metadata. To work around this issue, boot into rescue mode and run the command `dmraid -rE /dev/sdX` on the disks in question. Alternatively, run `dd if=/dev/zero of=/dev/sdX` and let it process up until the end of the disk. Note, however that this alternate procedure may take longer to complete and will erase all data on the disk.
- Installation of Red Hat Enterprise Linux 6 on an IBM ThinkPad T43 notebook may appear to stall after choosing storage options. In these circumstances, the installer is attempting to interact with the floppy drive, and may be unresponsive for up to 30 minutes.
- During the installation on POWER systems, the error messages similar to:

```
attempt to access beyond end of device
loop0: rw=0, want=248626, limit=248624
```

may be returned to `sys.log`. The errors do not prevent installation and only occur during initial setup. The filesystem created by the installer will function correctly.

- Installation on large disks (i.e. more than 2TB) on non-EFI platforms may encounter some limitations. Many BIOS systems can only boot disks that contain MSDOS partition tables, which

cannot fully address large disks. A GPT partition table can address the full disk, but may not be bootable from BIOS. Consequently, the Red Hat Enterprise Linux installer does not support installing the GRUB bootloader to disks that contain GPT partition tables on non-EFI systems. When installing Red Hat Enterprise Linux 6 on a non-EFI system that contains one or more large disks, create a GPT partition table on each of the disks before proceeding to the storage configuration portion of the install process. Leaving the large disks uninitialized, or using an MSDOS partition table on them, can cause problems when creating partitions using anaconda.

- Some Cisco UCS storage devices do not have UEFI support, which may lead to an unbootable Red Hat Enterprise Linux 6 system when installation is performed through virtual media with the system in "strict UCSM boot order rules" mode. Consequently, when installing using the UEFI method, after installation and reboot, the system will hang with a flashing cursor. To work around this issue, install the system using the BIOS install method as follows:
 1. Map the Red Hat Enterprise Linux 6 "boot.iso" file or entire OS DVD ISO using the virtual media tool
 2. Press F2 during boot to enter the BIOS setup screen
 3. Go to the "Boot Options" screen
 4. Change "UCSM boot order rules" to "Loose"
 5. Save settings and reboot
 6. Press F6 to access the boot device menu
 7. In the menu will be two options for the virtual media: "Cisco Virtual CD/DVD 1.20" and "EFI: Cisco Virtual CD/DVD 1.20 CDRom File1" select the first option to install using BIOS method. Note that only the first option will be present if using the "boot.iso" file, as it has no UEFI support.
 8. It may be necessary to re-order the devices in the BIOS Options screen after "Loose" mode has been selected in order to make the hard drive mapped to the system the first device in the boot order.

The use of BIOS install method will effectively work around the bug, but will prevent booting from disks using a GPT partition table. This will restrict the size of disks usable as a boot disk.

- When installing on the s390x architecture, if the installation is being performed over SSH, avoid resizing the terminal window containing the SSH session. If the terminal window is resized during installation, the installer will exit and installation will terminate.
- Multipath storage devices with serial numbers not exactly 16 or 32 characters in length will not be detected by anaconda during installation.
- Due to an issue with the shutdown sequence of the installer, Intel BIOS RAID sets might be left in an unclean condition post installation. Consequently, they will be rebuilt during the first boot of the system after installation. Note that this issue has no impact other than a slower first boot up after installation.
- The installer currently does not support having the /boot volume on a logical volume. Consequently, when setting up mount points during installation, the /boot volume cannot be on an LVM volume. System z supports /boot on an LVM volume. In order to exploit this, manual configuration after installation is required. Refer to the zipl documentation for further information.
- Minimal installations lack NetworkManager, so users wishing to have network interfaces configured for use on the first boot after installation need to make sure the network interfaces are configured

and the network service is enabled at boot time. The following kickstart commands will enable eth0 for DHCP and enable the 'network' service:

```
network --device eth0 --onboot yes --bootproto dhcp
services --enabled=network
```

Refer to the network device configuration documentation for more details on what the ifcfg-ethX files may contain.

- The kernel image provided on the CD/DVD is too large for Open Firmware. Consequently, on the POWER architecture, directly booting the kernel image over a network from the CD/DVD is not possible. Instead, use yaboot to boot from a network.
- The anaconda partition editing interface includes a button labeled **Resize**. Note that you can only shrink a partition with this button, not enlarge a partition.
- System z installations cannot use the ext4 filesystem for the boot partition. The recommended alternative filesystem is ext3.
- Channel IDs(read, write, data) for network devices are required for defining and configuring network devices on s390 systems. However, **system-config-kickstart** — the graphical user interface for generating a kickstart configuration — cannot define channel IDs for a network device. To work around this issue, manually edit the kickstart configuration that **system-config-kickstart** generates to include the desired network devices.
- During an MPATH installation on IBM POWER 7 systems, a "DiskLabelCommit Error" might be returned. To work around this issue, first install the system in a single path configuration. Connect to the system via SSH, clear the partitions using the **fdisk -l** command, and delete the partitions, then exit the SSH session. Finally, continue the installation from the installer.
- anaconda in Red Hat Enterprise Linux 6 for Power writes an incorrect value to /etc/rpm/macros that can cause issues when installing 32 and 64-bit PowerPC packages together. Users are advised to remove this file after installation.

2. Deployment

Upstart

In Red Hat Enterprise Linux 6, *init* from the *sysvinit* package has been replaced with *Upstart*, an event-based init system. This system handles the starting of tasks and services during boot, stopping them during shutdown and supervising them while the system is running. For more information on Upstart itself, refer to the **init(8)** man page.

Processes are known to Upstart as jobs and are defined by files in the **/etc/init** directory. Upstart is very well documented via man pages. Command overview is in **init(8)** and job syntax is described in **init(5)**.

Upstart provides the following behavioral changes in Red Hat Enterprise Linux 6:

- The **/etc/inittab** file is deprecated, and is now used *only* for setting up the default runlevel via the *initdefault* line. Other configuration is done via upstart jobs in the **/etc/init** directory.
- The number of active tty consoles is now set by the *ACTIVE_CONSOLES* variable in **/etc/sysconfig/init**, which is read by the **/etc/init/start-ttys.conf** job. The default value is *ACTIVE_CONSOLES=/dev/tty[1-6]*, which starts a getty on tty1 through tty6.

- A serial getty is still automatically configured if the serial console is the primary system console. In prior releases, this was done by **kudzu**, which would edit **/etc/inittab**. In Red Hat Enterprise Linux 6, configuration of the primary serial console is handled by **/etc/init/serial.conf**.
- To configure a getty running on a non-default serial console, you must now write an Upstart job instead of editing **/etc/inittab**. For example, if a getty on **ttyS1** is desired, the following job file (**/etc/init/serial-ttyS1.conf**) would work:

```
# This service maintains a getty on /dev/ttyS1.

start on stopped rc RUNLEVEL=[2345]
stop on starting runlevel [016]

respawn
exec /sbin/agetty /dev/ttyS1 115200 vt100-nav
```

As in prior releases, you should still make sure that **ttyS1** is in **/etc/securetty** if you wish to allow root logins on this getty.

There are some features from prior releases that are not supported in the move to Upstart. Among these are:

- Custom runlevels 7, 8 and 9. These custom runlevels can no longer be used.
- Using **/etc/shutdown.allow** for defining who can shut the machine down.

System z Performance

Some of the default tunables in Red Hat Enterprise Linux 6 are currently not optimally configured for System z workloads. Under most circumstances, System z machines will perform better using the following recommendations.

Dirty Ratio

It is recommended that the dirty ratio be set to 40 (Red Hat Enterprise Linux 6 default 20). Changing this tunable tells the system to not spend as much process time too early to write out dirty pages. Add the following line to **/etc/sysctl.conf** to set this tunable:

```
vm.dirty_ratio = 40
```

Scheduler

To increase the average time a process runs continuously and also improve the cache utilization and server style workload throughput at minor latency cost it is recommended to set the following higher values in **/etc/sysctl.conf**.

```
kernel.sched_min_granularity_ns = 10000000
kernel.sched_wakeup_granularity_ns = 15000000
kernel.sched_tunable_scaling = 0
kernel.sched_latency_ns = 80000000
```

Additionally, deactivating the Fair-Sleepers feature improves performance on a System z machine. To achieve this, set the following value in **/etc/sysctl.conf**

```
kernel.sched_features = 15834234
```

False positive hung task reports

It is recommended to prevent false positive hung task reports (which are rare, but might occur under very heavy overcommitment ratios). This feature can be used, but to improve performance, deactivate it by default by setting the following parameter in `/etc/sysctl.conf`:

```
kernel.hung_task_timeout_secs = 0
```

irqbalance service on the POWER architecture

On POWER architecture, the **irqbalance** service is recommended for automatic device Interrupt Request (IRQ) distribution across system CPUs to ensure optimal I/O performance. The **irqbalance** service is normally installed and configured to run during Red Hat Enterprise Linux 6 installation. However, under some circumstances, the **irqbalance** service is not installed by default. To confirm that the **irqbalance** service is running, execute the following command as root:

```
service irqbalance status
```

If the service is running, command will return a message similar to:

```
irqbalance (pid 1234) is running...
```

However, if the message lists the service as **stopped**, execute the following commands as root to start the **irqbalance** service:

```
service irqbalance start  
chkconfig --level 345 irqbalance on
```

If the output of the **service irqbalance status** command lists **irqbalance** as an **unrecognized service**, use **yum** to install the **irqbalance** package, and then start the service.

```
yum install irqbalance  
service irqbalance start
```



Note

The system does not need to be restarted after starting the **irqbalance** service

Setting the console log level

Use of the `LOGLEVEL` parameter in `/etc/sysconfig/init` to set the console loglevel is no longer supported. To set the console loglevel in Red Hat Enterprise Linux 6, pass `loglevel=<number>` as a boot time parameter.

Upgrading from previous pre-release versions

Upgrading to Red Hat Enterprise Linux 6 from Red Hat Enterprise Linux 5 or from previous pre-release versions of Red Hat Enterprise Linux 6 is not supported. If an upgrade of this type is attempted issues may be encountered including upgrading Java/OpenJDK packages. To work around this, manually remove the old packages and reinstall.

2.1. Known Issues

- When a system is configured to require smart card authentication, and there is no smartcard currently plugged into the system, then users might see the debug message:

```
ERROR: pam_pkcs11.c:334: no suitable token available'
```

This message can be safely ignored.

- Red Hat Enterprise Linux 6 Beta features Dovecot version 2.0. The configuration files used by Dovecot 2.0 are significantly different from those found in dovecot 1.0.x, the version shipped in previous releases of Red Hat Enterprise Linux. Specifically, `/etc/dovecot.conf` has been split into `/etc/dovecot/dovecot.conf` and `/etc/dovecot/conf.d/*.conf`
- Under some circumstances, the **readahead** service may cause the **auditd** service to stop. To work around this potential issue, disable the readahead collector by adding the following lines to the `/etc/sysconfig/readahead` configuration file:

```
READAHEAD_COLLECT="no"  
READAHEAD_COLLECT_ON_RPM="no"
```

Alternatively, the **readahead** package can be removed entirely.

- An error exists in the communication process between the samba daemon and the Common Unix Printing System (CUPS) scheduler. Consequently, the first time a print job is submitted to a Red Hat Enterprise Linux 6 system via Server Message Block (SMB), a timeout will occur. To work around this issue, use the following command to create a CUPS certificate before the first print job is submitted:

```
lpstat -E -s
```

- Under some circumstances, using the **rhnc_register** command to register a system with the Red Hat Network (RHN) might fail. When this issue is encountered, the `rhnc_register` command will return an error similar to:

```
# rhnc_register  
Segmentation fault (core dumped)  
or  
# rhnc_register  
***MEMORY-ERROR***: rhnc_register[11525]: GSlice: assertion failed:  
sinfo->n_allocated > 0  
Aborted (core dumped)
```

To work around this issue, set the following environment variable, then run the `rhnc_register` command again:

```
G_SLICE=always-malloc
```

- If a user has a `.bashrc` which outputs to `stderr`, the user will be unable to `sftp` into their account. From the user's point of view, the `sftp` session is immediately terminated after authentication.

2.1.1. Architecture Specific Known Issues

2.1.1.1. System z

The minimum hardware requirement to run Red Hat Enterprise Linux Beta is IBM System z9 (or better). The system may not IPL (i.e. boot) on earlier System Z hardware (e.g. z900 or z990)

2.1.1.2. IBM POWER (64-bit)

- When network booting an IBM POWER5 series system, you may encounter an error such as:

```
DEFAULT CATCH!, exception-handler=fff00300
```

If the path that locates the kernel and ramdisk is greater than 63 characters long, it will overflow a firmware buffer and the firmware will drop into the debugger.

POWER6 and POWER7 firmware includes a correction for this problem. Note that IBM POWER5 series is not a supported system.

- On some machines yaboot may not boot, returning the error message:

```
Cannot load ramdisk.image.gz: Claim failed for initrd memory at 02000000 rc=ffffffff
```

To work around this issue, change real-base from to **c00000**. Real-base can be obtained from OpenFirmware prompt with the **printenv** command and set with **setenv** command.

- Remote installs on IBM BladeCenter JS22 servers may encounter the following error message:

```
No video available. Your server may be in an unsupported resolution/refresh rate.
```

To work around this issue, specify the following GUI parameters:

```
video=SVIDEO-1:d radeon.svideo=0
```

- Some HP Proliant servers may report incorrect CPU frequency values in `/proc/cpuinfo` or `/sys/device/system/cpu/*/cpufreq`. This is due to the firmware manipulating the CPU frequency without providing any notification to the operating system. To avoid this ensure that the "HP Power Regulator" option in the BIOS is set to "OS Control". An alternative available on more recent systems is to set "Collaborative Power Control" to "Enabled".
- filecap crashes with a segmentation fault when run directly on an empty file. For example:

```
# filecap /path/to/empty_file
Segmentation fault (core dumped)
```

To work around this, run filecap on the directory that contains the empty file, and search the results for the required information. For example:

```
filecap /path/to/ | grep empty_file
```

- A change in the package that the sos tool uses to determine the installed version of Red Hat Enterprise Linux will cause the tool to incorrectly identify the major release version. This adversely impacts a small number of non-default sos plugins and may cause incomplete information to be captured from the system when these plugins are enabled. The affected plugins are:
 - general (only when using the non-default all_logs option)
 - cluster (diagnostics may not be run)

Users affected by this problem should retrieve any missing data manually from systems.

3. Virtualization

Para Virtualization on Hardware Virtualized Machines (PV on HVM)

Red Hat Enterprise Linux 6 guests under Red Hat Enterprise Linux 5 Xen hosts can now utilize the PV on HVM drivers to improve the performance of I/O on virtualized network devices (xen-vnif) and virtualized block storage devices.

To enable Xen PV on HVM support in a Red Hat Enterprise Linux 6 HVM guest, add the following to the kernel boot command line:

```
xen_pv_hvm=enable
```

Note, however, that due to conflicts with network configuration scripts, it is recommended that the xen guest vif specification set 'type=netfront' if the emulated rtl8139 device is not desired as the primary network interconnect.

virtio network device packet transmission algorithms

The virtio network device has two available algorithms for transmitting packets. The default is to use an asynchronous bottom half transmitter which typically shows good performance for all workloads. The alternate implementation uses a timer to delay transmit in an attempt to batch multiple packets together. The timer approach typically results higher latency, but may improve overall efficiency. To change from the default algorithm to the timer based approach, use the following procedure to create a wrapper script around qemu-kvm and specify it as the emulator for guests that require it.

1. create the wrapper script

```
$ cat > /usr/libexec/qemu-kvm.txtimer << EOF
#!/bin/sh
exec /usr/libexec/qemu-kvm `echo "$@" | sed
's|virtio-net-pci|virtio-net-pci,tx=timer|g'\`
EOF
```

2. Make script executable

```
$ chmod 755 /usr/libexec/qemu-kvm.txtimer
```

3. Set selinux permissions

```
$ restorecon /usr/libexec/qemu-kvm.txtimer
```

4. Create selinux module

```
$ cat > qemutxtimer.te << EOF
policy_module(qemutxtimer, 1.0)

gen_require(`
  attribute virt_domain;
  type qemu_exec_t;
  `)

can_exec(virt_domain, qemu_exec_t)
EOF
```

5. Build selinux module

```
$ make -f /usr/share/selinux/devel/Makefile
```

6. Install selinux module

```
$ semodule -i qemutxtimer.pp # May later be uninstalled with -r
```

7. Update guest XML to use qemu-kvm wrapper

```
$ virsh edit $GUEST
```

Replace:

```
<emulator>/usr/libexec/qemu-kvm</emulator>
```

With:

```
<emulator>/usr/libexec/qemu-kvm.txtimer</emulator>
```

3.1. Known Issues

- Under some circumstances, installation of a Red Hat Enterprise Linux 6 virtual guest stalls after the optional testing of media. Note that this issue has only been observed with Red Hat Enterprise Linux 6 guests that utilize multiple virtualized CPUs. To work around this issue, use a media source that is known to be verified, and skip the media test, or use a single virtualized CPU during installation.
- Cancelling the disk physical cache for block devices and use of barriers for filesystems may slow down qcow2 dramatically. Use the following command to reduce the frequency of sync requests by pre-allocating new images and setting the cluster size to 2M

```
./qemu-img create -opreallocation=metadata -ocluster_size=2M -f qcow2 $DISK $SIZE
```

- In earlier versions of Red Hat Enterprise Linux, libvirt permitted PCI devices to be insecurely assigned to guests. In Red Hat Enterprise Linux 6, assignment of insecure devices is disabled by default by libvirt. However, this may cause assignment of previously working devices to start failing.

To enable the old, insecure setting, edit `/etc/libvirt/qemu.conf`, set `"relaxed_acs_check = 1"`, and restart `libvirtd`. Note that this action will re-open possible security issues.

- Users upgrading from pre-release versions of Red Hat Enterprise Linux 6 (i.e. the **virt-v2v** versions less than `virt-v2v-0.6.2-2.el6`) may be required to update the default `virt-v2v` configuration file. Specifically, the `'viostor'` app for Windows guests is replaced by the `'virtio'` app, which now points to the directory containing the complete driver. Refer to the updated default configuration file for further details.
- I/O Advanced Programmable Interrupt Controller (I/O APIC) timer interrupts are not emulated as non-maskable interrupts (NMIs) to virtualized guests. Consequently, if a virtualized guest uses the kernel parameter `nmi_watchdog=1`, the guest kernel will panic on boot.
- The balloon service on Windows 7 guests can only be started by the "Administrator" user.
- Direct Asynchronous IO (AIO) that is not issued on filesystem block boundaries, and falls into a hole in a sparse file on ext4 or xfs filesystems, may corrupt file data if multiple I/O operations modify the same filesystem block. Specifically, if `qemu-kvm` is used with the `aio=native` IO mode over a sparse device image hosted on the ext4 or xfs filesystem, guest filesystem corruption will occur if partitions are not aligned with the host filesystem block size. Generally, do not use `aio=native` option along with `cache=none` for QEMU. This issue can be avoided by using one of the following techniques:
 1. Align AIOs on filesystem block boundaries, or do not write to sparse files using AIO on xfs or ext4 filesystems.
 2. KVM: Use a non-sparse system image file or allocate the space by zeroing out the entire file.
 3. KVM: Create the image using an ext3 host filesystem instead of ext4.
 4. KVM: Invoke `qemu-kvm` with `aio=threads` (this is the default).
 5. KVM: Align all partitions within the guest image to the host's filesystem block boundary (default 4k).
- On Red Hat Enterprise Linux 6 KVM virtual guests, unmounting a filesystem on an mdraid volume does not immediately free the underlying device for the `mdadm --stop operation`. Consequently, during installation on a system with pre-existing mdraid volumes the following error can appear while `anaconda` is looking for storage devices:

```
MDRaidError: mddeactivate failed for /dev/md1: 08:26:59,485 ERROR : Perhaps a running process, mounted filesystem or active volume group?
```

To work around this issue, erase all data on the volume before installation by clearing the first several sectors of the volume with zeros.

- Libvirt uses transient iptables rules for managing NAT or bridging to virtual machine guests. Any external command that reloads iptables state (such as running `system-config-firewall`) will overwrite the entries needed by libvirt. Consequently, after running any command or tool that changes the state of iptables, guests may lose access the network. To work around this issue, use the command `'service libvirt reload'` to restore libvirt's additional iptables rules.
- Adding an `rtl8139` NIC to an active Windows 2008 guest may result in the `qemu-kvm` process exiting. To work around this issue, shutdown the guest before adding additional `rtl8139` NICs. Alternatively, install the `virtio-net` drivers and add a `virtio` NIC.
- KVM users with a mix of `virtio` and `ata` disks should verify the boot device that `anaconda` chooses during installation. To verify the boot device, locate the "Install Target Devices" list in the disk

selection screen that follows the partitioning type screen. Verify the boot device selection, which is indicated by a selector in the left-most column of the "Install Target Devices" list.

- When installing Red Hat Enterprise Linux 6 as a new KVM guest, installer may incorrectly report amount of free memory available. Consequently, installation may terminate or switch to the text user interface. To work around this issue, increase amount of RAM allocated for the guest to 128 MB more than specified for the architecture and installation method.
- A Windows virtual machine must be restarted after the installation of the kernel windows driver framework. If the virtual machine is not restarted it may crash when a memory balloon operation is performed.
- Under some circumstances, if an 82576 Network driver (igb) is reloaded with the `max_vfs=8` parameter and an uncorrectable PCIe AER error is seen on its port, the operation will hang or crash the host system. This error has been encountered with two 82576 devices connected via an IDT PES12N3A PCI Express Switch (rev 0c) plugged into a Westmere-EP's 5520/5500/X58 I/O Hub PCI Express Root Port 3. Note that other 82576 devices and IDT switches have worked in other Westmere-based systems

If the error occurs, two workarounds have been found to enable the use of all eight virtual functions (VFs) for guest virtual machines (VMs):

1. Reload the 82576 driver with `max_vfs=1`, then unload, then reload with `max_vfs=8`. For example:

```
rmmod igb
modprobe igb max_vfs=1
rmmmod igbvf
rmmod igb
modprobe igb max_vfs=8
```

2. If PCI AER functionality is not needed in the host, boot the kernel with the parameter setting: `pci=noaer`

- A dual function, 82576 interface (codename: Kawela, PCI Vendor/Device ID: 8086:10c9) cannot have both physical functions (PF's) device-assigned to a Windows 2008 guest. Either physical function can be device assigned to a Windows 2008 guest (PCI function 0 or function 1), but not both.
- virt-v2v is able to convert guests running on ESX server. A current limitation in virt-v2v means that if an ESX guest has a disk with a snapshot, the snapshot must be on the same datastore as the underlying disk storage. If the snapshot and underlying storage are on different datastores, virt-v2v will report a 404 error while trying to retrieve the storage.
- Under some circumstances, the virtio queue will fill if an application on a guest repeatedly writes to a **virtio-serial** character device while the host is not processing the queue. Consequently, the guest will enter an infinite loop and appear to be hung. Once the host side of the character device is read from, the guest will return to normal functionality.
- The `qemu-kvm` options to enable VMware device emulation are not functional or supported in Red Hat Enterprise Linux 6.
- Avoid running `guestfish` (without the `--ro` option), `virt-edit`, `virt-tar` (in upload mode), `virt-win-reg` (in merge mode) or `guestmount` (without the `--ro` option) on live virtual machine disks. If any of these tools are used on live virtual machines, disk corruption might occur.

4. Storage and Filesystems

The ext4 Filesystem

The ext4 file system is a scalable extension of the ext3 file system, which was the default file system of Red Hat Enterprise Linux 5. Ext4 is now the default file system of Red Hat Enterprise Linux 6

Because of delayed allocation and other performance optimizations, ext4's behavior of writing files to disk is different from ext3. In ext4, a program's writes to the file system are not guaranteed to be on-disk unless the program issues an `fsync()` call afterwards.

Further information on the allocation features of ext4 is available in the [Storage Administration Guide](#)⁴⁹

CIFS servers that require plaintext passwords

Some Common Internet File System (CIFS) servers require plaintext passwords for authentication. Support for plaintext password authentication can be enabled using the command:

```
echo 0x37 > /proc/fs/cifs/SecurityFlags
```



Warning

This operation can expose passwords by removing password encryption.

Event Tracing in GFS2

GFS2's event tracing is provided via the generic tracing infrastructure. The events are designed to be useful for debugging purposes. Note, however that it is not guaranteed that the GFS2 events will remain the same throughout the lifetime of Red Hat Enterprise Linux 6. Further details on GFS2's gloocks and event tracing can be found in the following 2009 Linus Symposium paper: <http://kernel.org/doc/ols/2009/ols2009-pages-311-318.pdf>

mpi-selector

The `mpi-selector` package has been deprecated in Red Hat Enterprise Linux 6. **environment-modules** is now used to select which Message Passing Interface (MPI) implementation is to be used.



Note

The man page for the `module` command contains detailed documentation for the **environment-modules** package.

To return a list of what modules are available, use:

⁴⁹ http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/newfilesys-ext4.html

```
module avail
```

To load or unload a module use the following commands:

```
module load <module-name>
module unload <module-name>
```

To emulate the behavior of `mpi-selector`, the module load commands must be placed in the shell init script (e.g. `/.bashrc`) to load the modules every login.

4.1. Technology Previews

fsfreeze

Red Hat Enterprise Linux 6 includes **fsfreeze** as a Technology Preview. **fsfreeze** is a new command that halts access to a filesystem on disk. **fsfreeze** is designed to be used with hardware RAID devices, assisting in the creation of volume snapshots. Further details on **fsfreeze** are in the **fsfreeze(8)** man page.

DIF/DIX support

DIF/DIX, is a new addition to the SCSI Standard and a Technology Preview in Red Hat Enterprise Linux 6. DIF/DIX increases the size of the commonly used 512-byte disk block from 512 to 520 bytes, adding the Data Integrity Field (DIF). The DIF stores a checksum value for the data block that is calculated by the Host Bus Adapter (HBA) when a write occurs. The storage device then confirms the checksum on receive, and stores both the data and the checksum. Conversely, when a read occurs, the checksum can be checked by the storage device, and by the receiving HBA.

The DIF/DIX hardware checksum feature must only be used with applications that exclusively issue `O_DIRECT` I/O. These applications may use the raw block device, or the XFS file system in `O_DIRECT` mode. (XFS is the only filesystem that does not fall back to buffered IO when doing certain allocation operations.) Only applications designed for use with `O_DIRECT` I/O and DIF/DIX hardware should enable this feature. Red Hat Enterprise Linux 6 includes the Emulex LPFC driver version 8.3.5.17, introducing support for DIF/DIX. For more information, refer to the [Storage Administration Guide](#)⁵⁵

Filesystem in Userspace

Filesystem in Userspace (FUSE) allows for custom filesystems to be developed and run in userspace.

LVM Snapshots of Mirrors

The LVM snapshot feature provides the ability to create backup images of a logical volume at a particular instant without causing a service interruption. When a change is made to the original device (the origin) after a snapshot is taken, the snapshot feature makes a copy of the changed data area as it was prior to the change so that it can reconstruct the state of the device. Red Hat Enterprise Linux 6 introduces the ability to take a snapshot of a mirrored logical volume.

A known issue exists with this Technology Preview. I/O might hang if a device failure in the mirror is encountered. Note, that this issue is related to a failure of the mirror log device, and that no work around is currently known.

⁵⁵ http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html-single/Storage_Administration_Guide/index.html#id3654852

btrfs

Btrfs is under development as a file system capable of addressing and managing more files, larger files, and larger volumes than the ext2, ext3, and ext4 file systems. Btrfs is designed to make the file system tolerant of errors, and to facilitate the detection and repair of errors when they occur. It uses checksums to ensure the validity of data and metadata, and maintains snapshots of the file system that can be used for backup or repair. The btrfs Technology Preview is only available on the x86_64 architecture.



Btrfs is still experimental

Red Hat Enterprise Linux 6 Beta includes Btrfs as a technology preview to allow you to experiment with this file system. You should not choose Btrfs for partitions that will contain valuable data or that are essential for the operation of important systems.

LVM Application Programming Interface (API)

Red Hat Enterprise Linux 6 Beta features the new LVM application programming interface (API) as a Technology Preview. This API is used to query and control certain aspects of LVM.

FS-Cache

FS-Cache is a new feature in Red Hat Enterprise Linux 6 Beta that enables networked file systems (e.g. NFS) to have a persistent cache of data on the client machine.

4.2. Known Issues

- Enterprise-class storage should always be mounted with using the **-o nobARRIER** option.
- When an LVM mirror suffers a device failure, a two-stage recovery takes place. The first stage involves removing the failed devices. This can result in the mirror being reduced to a linear device. The second stage — if configured to do so by the administrator — is to attempt to replace any of the failed devices. Note, however, that there is no guarantee that the second stage will choose devices previously in-use by the mirror that had not been part of the failure if others are available.
- In Red Hat Enterprise Linux 5, infiniband support (specifically the **openib** start script and the **openib.conf** file) were supplied by the **openib** package. In Red Hat Enterprise Linux 6, the **openib** package is renamed to **rdma**. Additionally, the service has been renamed to **rdma** and the configuration file is now located in **/etc/rdma/rdma.conf**.
- The NFSv4 server in Red Hat Enterprise Linux 6 currently allows clients to mount using UDP and advertises NFSv4 over UDP with rpcbind. However, this configuration is not supported by Red Hat and violates the RFC 3530 standard.
- If a device-mapper-multipath device is still open, but all of the attached paths have been lost, the device is unable to create a new table with no paths. Consequently, the following unusual output may be returned from the **multipath -ll output** command:

```
mpatha (3600a59a0000c2fd0003079284c122fec) dm-0,
size=2.0G hwhandler='0'
|+- policy='round-robin 0' prio=0 status=enabled
| `- #:#:#:# - #:# failed faulty running
`+- policy='round-robin 0' prio=0 status=enabled
|- #:#:#:# - #:# failed faulty running
```

```
`- #:##:## - #:# failed faulty running
```

Output of this type indicates that there are no paths to the device. The erroneous lines in the output preceded by the string `#:##:##` will be removed in a future release.

- **ext2** and **ext3** filesystems do not use a **page_mkwrite** mechanism to intercept page faults. The quota subsystem can not account for this additional usage when writing to disk. Consequently, a user may exceed their disk block quota by issuing memory-mapped writes into a sparse region of a file. Note, also, that this is a longstanding behavior in the ext2 and ext3 filesystems.
- **Parted** in Red Hat Enterprise Linux 6 cannot handle Extended Address Volumes (EAV) Direct Access Storage Devices (DASD) that have greater than 65535 cylinders. Consequently, EAV DASD drives cannot be partitioned using parted and installation on EAV DASD drives will fail. To work around this issue, complete the installation on a non EAV DASD drive, then add the EAV device after installation using the tools provided in **s390-utils**.
- Systems that have an Emulex FC controller (with SLI-3 based firmware) installed may return a kernel panic during install. If the SAN disk is not required for installation, work around this issue by disconnecting the SAN connection from the Emulex FC controller. Note that this issue does not occur on SLI-4 based controllers. To determine the firmware interface of the adapter, run the command

```
cat /sys/class/scsi_host/host{n}/fwrev
```

- When multipath is configured to use `user_friendly_names`, it stores the binding between the `wwid` and the alias in `/etc/multipath/bindings`. When multipath creates devices in early bootup, (for example when the root filesystem is on a multipath device) it looks at `/etc/multipath/bindings` in the `initramfs`. When it creates devices during normal operation, it looks at `/etc/multipath/bindings` in the root filesystem. Currently, these two files aren't synced during `initramfs` creation. Because of this, there may be naming conflicts which keep new multipath devices from being created after bootup. To work around this, the bindings for the devices created by the `initramfs` must be copied into `/etc/multipath/bindings` after installation. The format of the bindings is:

```
<alias><space><wwid>
```

for example:

```
mpatha 3600d0230000000000e13955cc3757801
```

- Direct Asynchronous IO (AIO) that is not issued on filesystem block boundaries, and falls into a hole in a sparse file on ext4 or xfs filesystems, may corrupt file data if multiple I/O operations modify the same filesystem block. Specifically, if `qemu-kvm` is used with the `aio=native` IO mode over a sparse device image hosted on the ext4 or xfs filesystem, guest filesystem corruption will occur if partitions are not aligned with the host filesystem block size. Generally, do not use `aio=native` option along with `cache=none` for QEMU. This issue can be avoided by using one of the following techniques:
 1. Align AIOs on filesystem block boundaries, or do not write to sparse files using AIO on xfs or ext4 filesystems.
 2. KVM: Use a non-sparse system image file or allocate the space by zeroing out the entire file.
 3. KVM: Create the image using an ext3 host filesystem instead of ext4.

4. KVM: Invoke `qemu-kvm` with `aio=threads` (this is the default).
 5. KVM: Align all partitions within the guest image to the host's filesystem block boundary (default 4k).
- Mixing the iSCSI **discoveryd** mode and the normal discovery mode is not supported. When using **discoveryd** mode, **iscsid** will attempt to login from all iSCSI **ifaces** found in `/var/lib/iscsi/ifaces`. If the **iface** cannot log into the target this will fill the log with failure messages every `discoveryd_poll_inval` seconds. To prevent this, the **iface** can be deleted by running `"iscsiadm -m iface -o delete -I ifaceName"`.
 - A change in the 2.6.31 Linux kernel made the `net.ipv4.conf.default.rp_filter = 1` more strict in the I/O that is accepted. Consequently, in Red Hat Enterprise Linux 6, if there are multiple interfaces on the same subnet and I/O is sent to the one that is not the default route, the I/O will be dropped. Note that this applies to iSCSI iface binding when multiple interfaces are on the same subnet. To work around this, set the `net.ipv4.conf.default.rp_filter` parameter in `/etc/sysctl.conf` to 0 or 2, and reboot the machine.
 - Attempting to run multiple LVM commands in quick succession might cause a backlog of these commands. Consequently, some of the operations requested might time-out, and subsequently, fail.
 - dracut currently only supports one FiberChannel over Ethernet (FCoE) connection to be used to boot from the root device. Consequently, booting from a root device that spans multiple FCoE devices (e.g. using RAID, LVM or similar techniques) is not possible.
 - If an LVM volume requires physical volumes that are multipath or FCoE devices, the LVM volume will not automatically activate. To enable automatic LVM activation, create a udev rules file `/etc/udev/rules.d/64-autolvm.rules` with the following content:

```
SUBSYSTEM!="block", GOTO="lvm_end"
ACTION!="add|change", GOTO="lvm_end"
KERNEL=="dm-[0-9]*", ACTION=="add", GOTO="lvm_end"
ENV{ID_FS_TYPE}!="LVM*_member", GOTO="lvm_end"

PROGRAM==" /bin/sh -c 'for i in $sys/$devpath/holders/dm-[0-9]*; do [ -e $$i ] && exit 0;
done; exit 1;' ", \
    GOTO="lvm_end"

RUN+=" /bin/sh -c '/sbin/lvm vgscan; /sbin/lvm vgchange -a y'"

LABEL="lvm_end"
```

Note, however that this work around may impact system performance.

- The `fscontext=`, `defcontext=`, `rootcontext=` or `context=` mount options should not be used for remount operations. Using these options can cause the remount of a manually mounted volume to fail, returning errors such as:

```
mount: /dev/shm not mounted already, or bad option
```

5. Networking

NetworkManager

NetworkManager is enabled by default if it is installed. However, **NetworkManager** is only installed by default in the client use cases. **NetworkManager** is available to be installed for the server use cases, but is not included in the default installation.

5.1. Technology Previews

IPv6 support in IPVS

The IPv6 support in IPVS (IP Virtual server) is considered Technology Preview.

5.2. Known Issues

- If the **qeth** interface was previously configured using system-config-network **1.6.0.e16.2**, the "OPTIONS=" line needs to be manually added to `/etc/sysconfig/network-scripts/ifcfg-<interface>`.

After the configuration has been manually changed, activate the interface by either rebooting the system, or running the following commands:

```
# /sbin/znet_cio_free
# SUBSYSTEM="ccw" DEVPATH="bus/ccw/devices/<SUBCHANNEL 0>" /lib/udev/ccw_init
# ifup <interface>
```

- A known issue in the bnx2 driver prevents BCM5709S network adapters from performing a vmcore core dump over NFS.
- Intel 82575EB ethernet devices do not function in a 32 bit environment. To work around this issue, modify the kernel parameters to include the **intel_iommu=off** option.
- Running the **rds-ping** command may fail, returning the error:

```
bind() failed, errno: 99 (Cannot assign requested address).
```

Note, also that this error may occur even with `LOAD_RDS=yes` set in `/etc/rdma/rdma.conf`. To work around this issue, load the **rds-tcp** module.

- Running the command **rds-stress** on a client may result in the following error attempting to connect to the server:

```
connecting to <server IP address>:4000: No route to host
connect(<server IP address>) failed#
```

- When configuring a network interface manually, including static IP addresses and search domains, it is possible that a **search** entry will not be propagated to `/etc/resolv.conf`. Consequently, short host names that do not include the domain name will fail to resolve. To workaroud this issue, add a **search** entry manually to `/etc/resolv.conf`.
- Under some circumstances, the NetworkManager panel applet cannot determine if a user has permission to enable networking. Consequently, after logging into the desktop, the "Enable Networking" and "Enable Wireless" checkboxes may be disabled. To work around this, run the following command as root:

```
touch /usr/share/polkit-1/actions/org.freedesktop.NetworkManager.policy
```

Alternatively, WiFi can be enabled using the command:

```
nmcli nm wifi on
```

or disabled using the command:

```
nmcli nm wifi off
```

- Under some circumstances, the **netcf** command crashes, returning the error message:

```
Failed to initialize netcf  
error: unspecified error
```

To work around this issue, set the following value in `/etc/sysctl.conf`:

```
net.bridge.bridge-nf-call-iptables = 0
```

This issue presents when the **augeas** library (used by **netcf**) has trouble parsing one of the system config files that netcf needs to read or modify.

- The default value of the Emulex lpfc module parameter, `lpfc_use_msi`, was 2 (MSI-X) on Red Hat Enterprise Linux 5.4. In Red Hat Enterprise Linux 6 this default is now set to 0 (INTx). This change causes the driver behavior to stop using MSI-X interrupt mode and reverts to using non-msi (INTx) interrupt mode. This change in defaults addresses apparent regressions in some hardware platforms, introduced when the default lpfc driver value was previously changed from 0 to 2 (which made MSI-X the default behavior).

If the lpfc module is behaving erratically, work around this issue by setting the lpfc module parameter `lpfc_use_msi` to 2.

6. Clustering

6.1. Technology Previews

pacemaker

Pacemaker, a scalable high-availability cluster resource manager, is included in Red Hat Enterprise Linux 6 as a Technology Preview. Pacemaker is not fully integrated with the Red Hat cluster stack.

6.2. Known Issues

- Supplying an invalid version number in `cluster.conf` as a parameter to the `cman_tool` command will cause the cluster to stop processing information. To work around this issue, ensure that the version number used is valid.

- Under some circumstances, creating cluster mirrors with the '--nosync' option may cause I/O to become extremely slow. Note that this issue only effects I/O immediately after the creation of the mirror, and only when '--nosync' is used. To work around this issue, run the following command after the creating the mirror.

```
lvchange --refresh <VG>/<LV>
```

- luci will not function with Red Hat Enterprise Linux 5 clusters unless each cluster node has ricci version 0.12.2-14
- The sync state of an inactive LVM mirror cannot be determined. Consequently, the primary device of an LVM mirror can only be removed when the mirror is in-sync.
- If device-mapper-multipath is used, and the default path failure timeout value (`/sys/class/fc_remote_ports/rport-xxx/dev_loss_tmo`) is changed, that the timeout value will revert to the default value after a path fails, and later restored. Note that this issue will present the lpfc, qla2xxx, ibmfsc or fnic Fibre Channel drivers. To work around this issue the `dev_loss_tmo` value must be adjusted after each path fail/restore event.
- Generally, placing mirror legs on different physical devices improves data availability. The command **lvcreate --alloc anywhere** does not guarantee placement of data on different physical devices. Consequently, the use of this option is not recommended. If this option is used, the location of the data placement must be manually verified.
- The GFS2 fsck program, `fsck.gfs2`, currently assumes that the gfs2 file system is divided into evenly-spaced segments known as resource groups. This is always the case on file systems formatted by `mkfs.gfs2`. It will also be the case for most file systems created as GFS (`gfs1`) and converted to gfs2 format with `gfs2_convert`. However, if a GFS file system was resized (with `gfs_grow`) while it was in the GFS format, the resource groups might not be evenly spaced. If the resource groups are not evenly spaced, and the resource groups or the resource groups index (`rindev`) become damaged, `fsck.gfs2` might not function correctly.

There is currently no workaround for this issue. However, if the resource groups are not damaged, avoid this issue by copying the file system contents to a new device with evenly-spaced resource groups. Format the new device as gfs2 with `mkfs.gfs2`, and copy the contents from the old device to the new device. The new device will have evenly-spaced resource groups.

7. Authentication

7.1. Technology Previews

certmonger

The certmonger service aims to manage certificates on behalf of services running on client systems. It warns administrators when a certificate which it has been asked to watch is nearing the end of its validity period, and can be told to attempt to automatically obtain a new certificate when this happens. It supports certificates and private keys stored in either PEM or NSS database formats. It can interact with CAs running either IPA or certmaster, and is intended to be extensible to support other implementations.

ipa-client

IPA is an integrated solution to provide centrally managed Identity (machine,user, virtual machines, groups, authentication credentials). This package includes client-side functionality that when combined with a supported server can be used to provide features like kerberized sshd.

7.2. Known Issues

- Enabling user authentication against an LDAP server using `authconfig --enableldapauth` does not correctly set up the `/etc/nslcd.conf` configuration file. Consequently, LDAP users will be denied access to the system. To work around this issue, remove the line containing `pam_password md5` from the `/etc/nslcd.conf` file.
- The System Security Services Daemon (SSSD) currently supports following LDAP referrals on anonymous-bind LDAP connections only.
- The authentication configuration utility does not keep the 'Require smart card for login' check box set when Kerberos is also enabled. When the check box is checked and the configuration is saved with the 'Apply' button, the system will correctly require smart card for login. However, on the subsequent run of the authentication configuration utility the check box will be unchecked again and it is necessary to check it again to keep the option switched on.
- When attempting to perform PKINIT pre-authentication, if the client has more than one possible candidate certificate the client may fail to select the certificate and key to use. This usually occurs if certificate selection is configured to use the value of the keyUsage extension, or if any of the candidate certificates does not contain a `subjectAltName` extension. Consequently, the client attempts to perform pre-authentication using a different (usually password-based) mechanism.
- After installing certmonger, the system message bus daemon needs to be signaled to reload its configuration to allow the certmonger service to start properly. To work around this issue, send the `dbus-daemon` process a `SIGHUP` signal, or, alternatively, reboot the system.

8. Security

8.1. Technology Previews

OpenSCAP

OpenSCAP is a set of open source libraries that support the Security Content Automation Protocol (SCAP) standards from the National Institute of Standards and Technology (NIST). OpenSCAP supports the SCAP components:

- Common Vulnerabilities and Exposures (CVE)
- Common Platform Enumeration (CPE)
- Common Configuration Enumeration (CCE)
- Common Vulnerability Scoring System (CVSS)
- Open Vulnerability and Assessment Language (OVAL)
- Extensible Configuration Checklist Description Format (XCCDF)

Additionally, the `openSCAP` package includes an application to generate SCAP reports about system configuration. This package is considered a Technology Preview in Red Hat Enterprise Linux 6.

TPM

TPM hardware can create, store and use RSA keys securely (without ever being exposed in memory), verify a platform's software state using cryptographic hashes and more. The user space libraries, `trousers` and `tpm-tools` are considered a Technology Preview in this Red Hat Enterprise Linux 6.

9. Devices and Device Drivers

PCI Device Ordering

In Red Hat Enterprise Linux 6, the PCI device ordering is based on the PCI device enumeration. PCI device enumeration is based on the PCI enumeration algorithm (depth first then breadth) and is constant per system type. Additionally, once the devices are discovered, the module loading process is sequentialized, providing persistent naming of the interfaces.

9.1. Technology Previews

Brocade BFA Driver

The Brocade BFA driver is considered a Technology Preview feature in Red Hat Enterprise Linux 6. The BFA driver supports Brocade FibreChannel and FCoE mass storage adapters.

SR-IOV on the be2net driver

The SR-IOV functionality of the Emulex be2net driver is considered a Technology Preview in Red Hat Enterprise Linux 6.

9.2. Known Issues

- The **udev** daemon in Red Hat Enterprise 6 watches all devices for changes. If a change occurs, the device is rescanned for device information to be stored in the udev database.

The scanning process causes additional I/O to devices after they were changed by tools. **udev** can be told to exclude devices from being watched with a **udev** rule. A rule can be created by adding a new file **<myname>.rules** in **/etc/udev/rules.d** containing the following line:

```
ACTION=="add|change", SYMLINK=="disk/by-id/scsi-SATA_SAMSUNG_HD400LDS0AXJ1LL903246",
OPTIONS+="nowatch"
```

The SYMLINK should be replaced with any symlink path found in **/dev/disk/*** for the device in question.

This will prevent unexpected I/O on the device, after data was written directly to the device (not on the filesystem). However, it will also prevent device updates in the udev database, like filesystem labels, symbolic links in **/dev/disk/***, etc.

- Under some circumstances, the **bfa-firmware** package in Red Hat Enterprise Linux 6 may cause these devices to encounter a rare memory parity error. To work around this issue, to update to the newer firmware package, available directly from Brocade.
- Red Hat Enterprise Linux 6 only has support for the first revision of the UPEK Touchstrip fingerprint reader (USB ID 147e:2016). Attempting to use a second revision device may cause the fingerprint reader daemon to crash. The command

```
lsusb -v -d 147e:2016 | grep bcdDevice
```

will return the version of the device being used in an individual machine.

- The Emulex Fibre Channel/Fibre Channel-over-Ethernet (FCoE) driver in Red Hat Enterprise Linux 6 does not support DH-CHAP authentication. DH-CHAP authentication provides secure access between hosts and mass storage in Fibre-Channel and FCoE SANs in compliance with the FC-SP specification. Note, however that the Emulex driver (**lpfc**) does support DH-CHAP authentication

on Red Hat Enterprise Linux 5, from version 5.4. Future Red Hat Enterprise Linux 6 releases may include DH-CHAP authentication.

- Partial Offload iSCSI adapters do not work on Red Hat Enterprise Linux. Consequently, devices that use the `be2iscsi` driver cannot be used during installation.
- The `hpsa_allow_any` kernel option allows the `hpsa` driver to be used with older hardware that typically uses the `cciss` module by default. To use the `hpsa` driver with older hardware, set `hpsa_allow_any=1` and blacklist the `cciss` module. Note, however that this is an unsupported, non-default configuration.
- Platforms with BIOS/UEFI that are unaware of PCI-e SR-IOV capabilities may fail to enable virtual functions
- The recommended minimum HBA firmware revision for use with the `mpt2sas` driver is "Phase 5 firmware" (i.e. with version number in the form `05.xx.xx.xx`.) Note that following this recommendation is especially important on complex SAS configurations involving multiple SAS expanders.
- The persistent naming of devices that are dynamically discovered in a system is a large problem that exists both in and outside of `kdump`. Nominally, devices are detected in the same order, which leads to consistent naming. In cases where devices are not detected in the same order, device abstraction layers (e.g. LVM) make essentially resolve the issue, though the use of metadata stored on the devices to create consistency. In the rare cases where no such abstraction layer is in use, and renaming devices causes issues with `kdump`, it is recommended that devices be referred to by disk label or UUID in `kdump.conf`.
- The following issues and limitations may be encountered with the Broadcom `bnx2`, `bnx2x`, and `cnic` drivers
 - Support for only one VLAN per port
 - If deactivating the interface (i.e. the `ifdown` and `ifup` commands) the driver will need to be unloaded and reloaded to function correctly.

10. Kernel

Kdump Auto Enablement

Kdump is now enabled by default on systems with large amounts of memory. Specifically, `kdump` is enabled by default on:

- systems with more than 4GB of memory on architectures with a 4KB page size (i.e. x86 or x86_64), or
- systems with more than 8GB of memory on architectures with larger than a 4KB page size (i.e. PPC64).

On systems with less than the above memory configurations, `kdump` is not auto enabled. Refer to `/usr/share/doc/kexec-tools-2.0.0/kexec-kdump-howto.txt` for instructions on enabling `kdump` on these systems.

`crashkernel` parameter syntax

Please note that in future versions of Red Hat Enterprise Linux 6 (i.e. Red Hat Enterprise Linux 6.1 and later) the `auto` value setting of the `crashkernel=` parameter (i.e. `crashkernel=auto`) will be deprecated.

Barrier Implementation in the Kernel

The barrier implementation in the Red Hat Enterprise Linux 6 kernel works by completely draining the I/O scheduler's queue, then issuing a preflush, a barrier, and finally a postflush request. However, since the supported file systems in Red Hat Enterprise Linux 6 all implement their own ordering guarantees, the block layer need only provide a mechanism to ensure that a barrier request is ordered with respect to other I/O already in the disk cache. This mechanism avoids I/O stalls experienced by queue draining. The block layer will be updated in future kernels to provide this more efficient mechanism of ensuring ordering.

Workloads that include heavy fsync or metadata activity will see an overall improvement in disk performance. Users taking advantage of the proportional weight I/O controller will also see a boost in performance. In preparation for the block layer updates, third party file system developers need to ensure that data ordering surrounding journal commits are handled within the file system itself, since the block layer will no longer provide this functionality.

These future block layer improvements will change some kernel interfaces such that symbols which are not on the kABI whitelist shall be modified. This may result in the need to recompile third party file system or storage drivers.

Systemtap Tracepoints

The following 3 virtual memory tracepoints are deprecated in Red Hat Enterprise Linux 6

- `trace_mm_background_writeout(unsigned long written)`
- `trace_mm_olddata_writeout(unsigned long written)`
- `trace_mm_balancedirty_writeout(unsigned long written)`

10.1. Technology Previews

Remote Audit Logging

The audit package contains the user space utilities for storing and searching the audit records generated by the audit subsystem in the Linux 2.6 kernel. Within the `audispd-plugins` subpackage is a utility that allows for the transmission of audit events to a remote aggregating machine. This remote audit logging application, `audisp-remote`, is considered a Technology Preview in Red Hat Enterprise Linux 6.

Linux (NameSpace) Container [LXC]

Linux (NameSpace) Containers [LXC] is a Technology Preview feature in Red Hat Enterprise Linux 6 Beta that provides isolation of resources assigned to one or more processes. A process is assigned a separate user permission, networking, filesystem name space from its parent.

Error Detection And Correction (EDAC) driver interface

The Error Detection And Correction (EDAC) driver interface for processors based on the Intel microarchitecture codename Nehalem is considered a Technology Preview in this pre-release version of Red Hat Enterprise Linux 6.

10.2. Known Issues

- Calgary IOMMU default detection has been disabled in this release. If you require Calgary IOMMU support add `'iommu=calgary'` as a boot parameter.
- The `kdump` service fails on systems with large amounts of memory and `crashkernel=auto` enabled, returning the error message `kdump: kexec: failed to load kdump kernel` in `/var/log/messages`.

To workaround this issue, change the `crashkernel` parameter to **128M** (on x86_64 and x86 architectures) or **256M** (on the ppc64 architecture).

- If the kdump crash recovery technology is enabled and in use on a given system, minimum memory requirements should be raised by the amount of memory reserved for kdump usage. This value is determined by the user, and specified on the kernel command line, via the `crashkernel` parameter. The default value for this setting is 128MB.

- When using the DIF/DIX hardware checksum features of a storage path behind a block device, errors will occur if the block device is used as a general purpose block device.

Buffered I/O or `mmap(2)` based IO will not work reliably as there are no interlocks in the buffered write path to prevent overwriting cached data while the hardware is performing DMA operations. An overwrite during a DMA operation will cause a torn write and the write will fail checksums in the hardware storage path. This problem is common to all block device or file system based buffered or `mmap(2)` I/O, so the problem of I/O errors during overwrites cannot be worked around.

DIF/DIX enabled block devices should only be used with applications that use `O_DIRECT` I/O. Applications should use the raw block device, though it should be safe to use the XFS file system on a DIF/DIX enabled block device if only `O_DIRECT` I/O is issued through the file system. In both cases the responsibility for preventing torn writes lies with the application, so only applications designed for use with `O_DIRECT` I/O and DIF/DIX hardware should enable this feature.

- The memory controller in Red Hat Enterprise Linux 6 beta may encounter stability issues when under heavy stress testing or memory pressure.
- The i686 debug kernel may crash on some systems when starting the udev service.
- Systems configured with Intel 82578DM NICs may not be recognized during boot/install resulting in driver load failure, (driver probe fails with error -2).
- This pre-release version of Red Hat Enterprise Linux 6 provides automated Physical CPU Socket and Memory Hot-Add support. Note, however, that CPU Socket and Memory Hot-Remove actions are not supported. Additionally, only single CPU Socket add events are supported at this time, and tsc support is disabled after a CPU Socket add event.
- In Beta releases of Red Hat Enterprise Linux 6, PCIe ASPM would be enabled on PCIe hierarchies even if they lacked an `_OSC` method as defined in section 4.5 of the PCI firmware specification, release 3.0. Post Beta, firmware must provide an appropriate `_OSC` method on all PCI roots in order to allow PCIe ASPM to be enabled. The `"pcie_aspm=force"` boot parameter may be passed in order to enable PCIe ASPM.
- Use of the `cciss` and `hpsa` drivers with some controllers (e.g. P400, P400i, E500, P800, P700m and 6402/6404) may cause kdump to fail.
- The top-level makefile to of the kernel in Red Hat Enterprise Linux 6 includes the `-Werror` option as part of the standard kernel build. Consequently, all kernel compile warnings are reported as errors. In non-production environments, the `-Werror` flag can be disabled by removing the following two lines from the top-level kernel Makefile:

```
KBUILD_CFLAGS += $(shell if [ $(CPP_VERS) -ge 4004004 ]; then \  
    echo "-Wno-array-bounds -Werror"; else echo ""; fi)
```

Note, however, that Red Hat does not support custom built kernels or custom built modules.

- Some SystemTap probes require the additional module, **uprobes.ko** at run time. This additional module is usually built automatically when the script is compiled. However, in the client-server case, the uprobes.ko module is not returned by the server to the client. Consequently, missing symbols are reported when the module representing the script is loaded. To work around this issue, use the following command to manually build the uprobes.ko module on the client host.

```
make -C <prefix>/share/systemtap/runtime/uprobes
```

Note that "<prefix>" is the install prefix for systemtap, and that this manual build of uprobes.ko will only need to be done once.

- Due to the way ftrace works when modifying the code during startup, the NMI watchdog causes too much noise and ftrace can not find a quiet period to instrument the code. Consequently, machines with more than 512 cpus will encounter issues with the NMI watchdog. Such issues will return error messages similar to "BUG: NMI Watchdog detected LOCKUP" and have either 'ftrace_modify_code' or 'ipi_handler' in the backtrace. To work around this issue, disable nmi_watchdog using the command:

```
nmi_watchdog=0
```

- Under some circumstances, a kernel panic on installation or boot may occur if the "Interrupt Remapping" feature is enabled in the BIOS. To work around this issue, disable interrupt remapping in the BIOS.
- The kernel will panic when booting the kdump kernel on a s390 system with an initramfs that contains an odd number of bytes. To work around this issue, generate an initramfs with sufficient padding such that it contains an even number of bytes.
- Creating many 'cpu' control groups (cgroups) on a system with a large number of CPUs will slow down the machine when the control groups feature is enabled. To work around this issue, disable control groups.
- Under certain circumstances, the Linux kernel makes an erroneous assumption about where to reserve memory for the kdump kernel on large-memory POWER systems. Consequently, a newly installed POWER system may return the following message during the initial post installation bootup:

```
returning from prom_init
Kernel panic - not syncing: ERROR: Failed to allocate 0x4000 bytes below 0x10000000.
Rebooting in 180 seconds..
```

Complete the following steps to work around this issue. Note, however, that this work around disables the kdump feature.

1. The system will reboot 180 seconds after the initial error message was returned. After reboot, the yaboot prompt will be presented:

```
Welcome to Red Hat Enterprise Linux!
Hit <TAB> for boot options
Welcome to yaboot version 1.3.14 (Red Hat 1.3.14-34.el6)
Enter "help" to get some basic usage information
boot:
```

At the prompt, enter the following line and press enter.

```
linux crashkernel=512M-2G:256M
```

2. Log in to the system as root, and open `/etc/yaboot.conf` in a text editor. The `yaboot.conf` file should be similar to:

```
# yaboot.conf generated by anaconda

boot=/dev/sda1
init-message="Welcome to Red Hat Enterprise Linux!\nHit <TAB> for boot options"

partition=2
timeout=5
install=/usr/lib/yaboot/yaboot
delay=30
enablecdboot
enableofboot
enablenetboot
nonvram
fstype=raw

image=/vmlinuz-2.6.32-59.el6.ppc64
    label=linux
    read-only
    initrd=/initramfs-2.6.32-59.el6.ppc64.img
    append="rd_NO_LUKS rd_NO_LVM rd_NO_MD rd_NO_DM LANG=en_US.UTF-8
SYSFONT=latarcyrheb-sun16 KEYTABLE=us console=hvc0 crashkernel=auto rhgb quiet
root=UUID=63f94acf-6241-4a66-a861-9de912602287"
```

Remove the string `crashkernel=auto` from the `append=` line. Save the file, and exit the editor. Subsequent reboots of the system will boot to the system prompt.

- On 64-bit POWER systems the EHEA NIC driver will fail when attempting to dump a vmcore via NFS. To work around this issue, utilize other kdump facilities, for example dumping to the local filesystem, or dumping over SSH.
- A BIOS emulated floppy disk might cause the installation or kernel boot process to hang. To avoid this, disable emulated floppy disk support in the BIOS.
- The preferred method to enable `nmi_watchdog` on 32-bit x86 systems is to use either `nmi_watchdog=2` or `nmi_watchdog=lapic` parameters. The parameter `nmi_watchdog=1` is not supported.
- The module loading operation of certain crypto libraries will not be successful. Consequently, the modules required for *in-kernel crypto* cannot be loaded. **In-kernel crypto** cannot be used with Red Hat Enterprise Linux 6 until this issue is resolved.
- A BIOS issue on some platforms incorrectly indicates that the system busmastering flag must be checked before entering the deep C state. Consequently, some systems might spend a significantly lower percentage of time in deep C states (C3 and lower) in Red Hat Enterprise Linux 6 compared to Red Hat Enterprise Linux 5.5. Updated the BIOS on affected systems will resolve this issue.
- IMA in Red Hat Enterprise Linux 6.0 GA is enabled by loading an IMA policy. However, future updates will require the boot parameter `"ima=on"` in addition to loading an IMA policy to enable IMA. This change reduces overhead on systems not using IMA.

11. Development and Tools

11.1. Technology Previews

libdfp

An updated libdfp library is available in Red Hat Enterprise Linux 6. libdfp is a decimal floating point math library, and is available as an alternative to the glibc math functions on Power and s390x architectures, and is available in the supplementary channels.

Eclipse Plugins

The following plugins for the Eclipse software development environment are considered to be Technology Previews in this pre-release version of Red Hat Enterprise Linux 6

- The Mylyn plugin for the Eclipse task management subsystem
- the **eclipse-callgraph** C/C++ Call Graph Visualization plugin

11.2. Known Issues

- cURL is a tool for getting files from FTP, HTTP, Gopher, Telnet, and Dict servers, using any of the supported protocols. The cURL API, and consequently, the python bindings for cURL, do not provide textual messages for errors. Therefore, all applications that use the python bindings for cURL will return errors in formats such as:

```
Pycurl Error 6 - ""
```

instead of more useful messages such as:

```
Pycurl Error 6 - "Could not resolve hostname: blah.example.com"
```

cURL error codes can be manually interpreted by reading the `/usr/include/curl/curl.h` file.

- Due to a deficiency in java-1.6.0-ibm-plugin for AMD64 and Intel 64, IBM Java 6 Web Start cannot open JNLP files. This affects file management tools and WWW browsers. To work around this open JNLP files using the command:

```
/usr/lib/jvm/jre-1.6.0-ibm.x86_64/bin/javaws file.jnlp
```

Note that 32-bit packages are not affected by this issue.

- Under some circumstances on the PPC64 architecture, Ruby does not save the context correctly when switching threads. Consequently, when a thread is restored it has a stale value which might return a architecture fault.
- Under some circumstances, libdfp encounters an issue converting some values from string to DFP with the conversion command `strtod32`. The `strtod64` and `strtod128` commands do work correctly.

12. Desktop

nautilus-open-terminal behavior change

The **nautilus-open-terminal** package provides a right-click "Open Terminal" option to open a new terminal window in the current directory. Previously, when this option was chosen from the Desktop, the new terminal window location defaulted to the user's home directory. However, in Red Hat Enterprise Linux 6, the default behavior opens the Desktop directory (i.e. `~/Desktop/`). To enable the previous behavior, use the following command to set the **desktop_opens_home_dir** GConf boolean to true:

```
gconftool-2 -s /apps/nautilus-open-terminal/desktop_opens_home_dir --type=bool true
```

Adobe Flash and Adobe Acrobat Reader on 64-Bit

The 64-bit Red Hat Enterprise Linux Supplementary CD contains the 32-bit versions of Adobe Acrobat Reader and Adobe Flash for use on the 64-bit architecture. To use these browser plugins correctly, the **nspluginwrapper.i686** and **alsa-plugins-pulseaudio.i686** packages must be installed prior to the installation of the plugins.

gnome-packagekit architecture filter

By default, gnome-packagekit uses a filter to hide packages that are not the same architecture as the system. Consequently, when installing packages for other architectures (e.g. the 32-bit versions of `acroread` and `flash-plugin` on the 64-bit architecture) the "Only native filters" from the Filters menu must be unchecked for these packages to be visible.

12.1. Known Issues

- When enabled, fingerprint authentication is the default authentication method to unlock a workstation, even if the fingerprint reader device is not accessible. However, after a 30 second wait, password authentication will become available.
- ATI RN50/ES1000 graphics devices have limited Video RAM (VRAM) and are restricted to an 8-bit color depth for the text console. Consequently, the graphical boot screen is unavailable on systems using these graphics devices.
- On the GNOME desktop, the CD/DVD burning utility `brasero` conflicts with the automounting feature in Nautilus. Consequently, the following error message will be displayed when `brasero` attempts to verify the checksum of the disc:

```
Error while burning: You do not have the required permissions to use this drive
```

In most cases, the data is still written to the disc.

- The **system-config-users** tool cannot always detect if a home directory can be created correctly. Consequently, `system-config-users` might fail silently when attempting to create a home directory on some file systems (e.g. home directories located beneath an `autofs` mount-point). Typically, when this issue is encountered, the user account itself is created, but the creation of the home directory fails. To create a user with an auto-mounted home directory, create the home directory manually before creating the user in `system-config-users`.
- Evolution's IMAP backend only refreshes folder contents under the following circumstances: when the user switches into or out of a folder, when the auto-refresh period expires, or when the user manually refreshes a folder (i.e. using the menu item **Folder > Refresh**). Consequently, when

replying to a message in the Sent folder, the new message does not immediately appear in the Sent folder. To see the message, force a refresh using one of the methods describe above.

- Not all languages have predefined default input method engines. Consequently, in some languages, **ibus** will not have an input method engine configured. To work around this issue, add an input method using the Input Method configuration dialog (**System > Preferences > Input Method**
- Using the im-chooser tool, XIM cannot be disabled as the default GTK immodule. Disabling input-methods using im-chooser and restarting the desktop session will still result in GTK applications using the XIM immodule. Consequently, using the Ctrl+Shift+U key combination to the directly input of Unicode characters from their hexadecimal code will not work. To work around this issue, use im-chooser to enable ibus. Enabling ibus permits gtk-im-context-simple's Unicode input and compose sequences to be used.
- The hardware mute button on Lenovo ThinkPad X200 notebooks does not work. Note, however, that the volume down and volume up buttons function correctly.
- The clock applet in the GNOME panel has a default location of Boston, USA. Additional locations are added by via the applet's preferences dialog. Additionally, to change the default location, left-click the applet, hover over the desired location in the "Locations" section, and click the "Set..." button that appears.
- In some multi-monitor configurations (e.g. dual monitors with both rotated), the cursor confinement code produces incorrect results. For example, the cursor may be permitted to disappear offscreen when it should not, or be prevented from entering some areas where it should be allowed to go. Currently, the only work around to this issue is to disable monitor rotation.
- ATI RN50/ES1000 graphics devices have a lower number of hardware controllers than output connectors. Due to a defect in the graphical boot system, this type of configuration results in a blank display. Consequently, users of systems with these ATI graphics devices will experience prolonged (potentially up to 2 minutes or longer) blank screens during boot up and shutdown. Once the boot process completes and a login prompt is available, the display will function as expected. The prolonged blank screen can be avoided by removing "rhgb" from the list of boot parameters on the kernel command line in **/etc/grub.conf**
- If a Russian keyboard is chosen during system installation, the login screen is configured to use Russian input for user names and passwords by default. However, pressing Left Shift and Right Shift does not cause the input to change to ASCII mode. Consequently, the user cannot log in. To work around this issue, run the following sequence, as root, post installation:

```
. /etc/sysconfig/keyboard; echo $LAYOUT | grep -q ",us" && gconftool-2
--direct --config-source xml:readwrite:/var/lib/gdm/.gconf --set
/apps/gdm/simple-greeter/recent-layouts --type list --list-type string $(echo
$LAYOUT | awk -F, '{ print "[" $2 ", " $1 "]"; }') && echo "DONE"
```

- For KMS drivers, the syntax is:

```
video=[connector:]mode
```

"connector", which is optional maps to the name of the connector as listed in `/sys/class/drm/card0`. For example:

```
~% ls /sys/class/drm/card0
```

```
card0-LVDS-1 card0-VGA-1 dev device power subsystem uevent
```

This device has connectors named LVDS-1 and VGA-1. If no connector is specified the requested mode will apply to all connectors.

Mode strings may be of the form:

```
<xres>x<yres>[R][-<bpp>][@<refresh>][i][eDd]
```

Parts inside <> are mandatory, parts inside [] are optional. R requests the use of the CVT reduced-blanking formula, applicable for some digital displays; otherwise GTF is used. i requests an interlaced mode. e forces the output to be enabled even if it appears to be disconnected; d forces the output to be disabled. For DVI connections, D forces the use of the digital signal path instead of analog; on other connectors it has no effect. Only one of e, d, or D may be given.

- Under some circumstances, the Add/Remove Software (gpk-application) graphical user interface does not display Supplementary groups or packages the Supplementary group is chosen. To work around this, use the System>Refresh Package Lists option to refresh the package lists.

A. Package Manifest

Previous versions of the Technical Notes contained a Package Manifest appendix. The [Package Manifest is now available as a separate document](#)¹.

B. Large Package Updates

B.1. kernel

B.1.1. RHSA-2011:0883 - Important: kernel security and bug fix update



Important

This update has already been released as the security errata [RHSA-2011:0883](#)¹

Updated kernel packages that fix multiple security issues and three bugs are now available for Red Hat Enterprise Linux 6.0 Extended Update Support.

The Red Hat Security Response Team has rated this update as having important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links after each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

¹ http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Package_Manifest/

¹ <https://rhn.redhat.com/errata/RHSA-2011-0883.html>

This update includes backported fixes for security issues. These issues, except for [CVE-2011-1182](https://www.redhat.com/security/data/cve/CVE-2011-1182.html)², only affected users of Red Hat Enterprise Linux 6.0 Extended Update Support as they have already been addressed for users of Red Hat Enterprise Linux 6 in the 6.1 update, [RHSA-2011:0542](https://rhn.redhat.com/errata/RHSA-2011-0542.html)³.

Security fixes:

- * Buffer overflow flaws were found in the Linux kernel's Management Module Support for Message Passing Technology (MPT) based controllers. A local, unprivileged user could use these flaws to cause a denial of service, an information leak, or escalate their privileges. ([CVE-2011-1494](https://www.redhat.com/security/data/cve/CVE-2011-1494.html)⁴, [CVE-2011-1495](https://www.redhat.com/security/data/cve/CVE-2011-1495.html)⁵, Important)
- * A flaw was found in the Linux kernel's networking subsystem. If the number of packets received exceeded the receiver's buffer limit, they were queued in a backlog, consuming memory, instead of being discarded. A remote attacker could abuse this flaw to cause a denial of service (out-of-memory condition). ([CVE-2010-4251](https://www.redhat.com/security/data/cve/CVE-2010-4251.html)⁶, [CVE-2010-4805](https://www.redhat.com/security/data/cve/CVE-2010-4805.html)⁷, Moderate)
- * A flaw was found in the Linux kernel's Transparent Huge Pages (THP) implementation. A local, unprivileged user could abuse this flaw to allow the user stack (when it is using huge pages) to grow and cause a denial of service. ([CVE-2011-0999](https://www.redhat.com/security/data/cve/CVE-2011-0999.html)⁸, Moderate)
- * A flaw in the Linux kernel's Event Poll (epoll) implementation could allow a local, unprivileged user to cause a denial of service. ([CVE-2011-1082](https://www.redhat.com/security/data/cve/CVE-2011-1082.html)⁹, Moderate)
- * An inconsistency was found in the interaction between the Linux kernel's method for allocating NFSv4 (Network File System version 4) ACL data and the method by which it was freed. This inconsistency led to a kernel panic which could be triggered by a local, unprivileged user with files owned by said user on an NFSv4 share. ([CVE-2011-1090](https://www.redhat.com/security/data/cve/CVE-2011-1090.html)¹⁰, Moderate)
- * It was found that some structure padding and reserved fields in certain data structures in KVM (Kernel-based Virtual Machine) were not initialized properly before being copied to user-space. A privileged host user with access to /dev/kvm could use this flaw to leak kernel stack memory to user-space. ([CVE-2010-3881](https://www.redhat.com/security/data/cve/CVE-2010-3881.html)¹¹, Low)
- * A missing validation check was found in the Linux kernel's `mac_partition()` implementation, used for supporting file systems created on Mac OS operating systems. A local attacker could use this flaw to cause a denial of service by mounting a disk that contains specially-crafted partitions. ([CVE-2011-1010](https://www.redhat.com/security/data/cve/CVE-2011-1010.html)¹², Low)
- * A buffer overflow flaw in the DEC Alpha OSF partition implementation in the Linux kernel could allow a local attacker to cause an information leak by mounting a disk that contains specially-crafted partition tables. ([CVE-2011-1163](https://www.redhat.com/security/data/cve/CVE-2011-1163.html)¹³, Low)
- * Missing validations of null-terminated string data structure elements in the `do_replace()`, `compat_do_replace()`, `do_ip6t_get_ctl()`, `do_ip6t_get_ctl()`, and `do_arpt_get_ctl()`

² <https://www.redhat.com/security/data/cve/CVE-2011-1182.html>

³ <https://rhn.redhat.com/errata/RHSA-2011-0542.html>

⁴ <https://www.redhat.com/security/data/cve/CVE-2011-1494.html>

⁵ <https://www.redhat.com/security/data/cve/CVE-2011-1495.html>

⁶ <https://www.redhat.com/security/data/cve/CVE-2010-4251.html>

⁷ <https://www.redhat.com/security/data/cve/CVE-2010-4805.html>

⁸ <https://www.redhat.com/security/data/cve/CVE-2011-0999.html>

⁹ <https://www.redhat.com/security/data/cve/CVE-2011-1082.html>

¹⁰ <https://www.redhat.com/security/data/cve/CVE-2011-1090.html>

¹¹ <https://www.redhat.com/security/data/cve/CVE-2010-3881.html>

¹² <https://www.redhat.com/security/data/cve/CVE-2011-1010.html>

¹³ <https://www.redhat.com/security/data/cve/CVE-2011-1163.html>

functions could allow a local user who has the **CAP_NET_ADMIN** capability to cause an information leak. ([CVE-2011-1170](#)¹⁴, [CVE-2011-1171](#)¹⁵, [CVE-2011-1172](#)¹⁶, Low)

* A missing validation check was found in the Linux kernel's signals implementation. A local, unprivileged user could use this flaw to send signals via the `sigqueueinfo` system call, with the **si_code** set to `SI_TKILL` and with spoofed process and user IDs, to other processes. Note: This flaw does not allow existing permission checks to be bypassed; signals can only be sent if your privileges allow you to already do so. ([CVE-2011-1182](#)¹⁷, Low)

Red Hat would like to thank Dan Rosenberg for reporting CVE-2011-1494 and CVE-2011-1495; Nelson Elhage for reporting CVE-2011-1082; Vasily Kulikov for reporting CVE-2010-3881, CVE-2011-1170, CVE-2011-1171, and CVE-2011-1172; Timo Warns for reporting CVE-2011-1010 and CVE-2011-1163; and Julien Tinnes of the Google Security Team for reporting CVE-2011-1182.

Bug fixes:

BZ#[590187](#)¹⁸

Previously, CPUs kept continuously locking up in the `inet_csk_bind_conflict()` function until the entire system became unreachable when all the CPUs were unresponsive due to a hash locking issue when using port redirection in the `__inet_inherit_port()` function. With this update, the underlying source code of the `__inet_inherit_port()` function has been modified to address this issue, and CPUs no longer lock up.

BZ#[709380](#)¹⁹

A previously released patch for BZ#[625487](#)²⁰ introduced a kABI (Kernel Application Binary Interface) workaround that extended **struct sock** (the network layer representation of sockets) by putting the extension structure in the memory right after the original structure. As a result, the **prot->obj_size** pointer had to be adjusted in the `proto_register` function. Prior to this update, the adjustment was done only if the `alloc_slab` parameter of the `proto_register` function was not `0`. When the `alloc_slab` parameter was `0`, drivers performed allocations themselves using `sk_alloc` and as the allocated memory was lower than needed, a memory corruption could occur. With this update, the underlying source code has been modified to address this issue, and a memory corruption no longer occurs.

BZ#[706543](#)²¹

An **IDX_ACTIVATE** timeout occurred during an online setting of an OSN device. This was because an incorrect function was provided on the **IDX_ACTIVATE**. Because OSN devices use the same function level as OSD devices, this update adds OSN devices to the initialization function for the **func_level**; thus, resolving this issue.

Users should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

¹⁴ <https://www.redhat.com/security/data/cve/CVE-2011-1170.html>

¹⁵ <https://www.redhat.com/security/data/cve/CVE-2011-1171.html>

¹⁶ <https://www.redhat.com/security/data/cve/CVE-2011-1172.html>

¹⁷ <https://www.redhat.com/security/data/cve/CVE-2011-1182.html>

¹⁸ https://bugzilla.redhat.com/show_bug.cgi?id=590187

¹⁹ https://bugzilla.redhat.com/show_bug.cgi?id=709380

²⁰ https://bugzilla.redhat.com/show_bug.cgi?id=625487

²¹ https://bugzilla.redhat.com/show_bug.cgi?id=706543

B.1.2. RHSA-2011:0498 - Important: kernel security, bug fix and enhancement update



Important

This update has already been released as the security errata [RHSA-2011:0498](#)²²

Updated kernel packages that resolve several security issues, fix various bugs and add an enhancement are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links after each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security fixes:

- * An integer overflow flaw in `ib_uverbs_poll_cq()` could allow a local, unprivileged user to cause a denial of service or escalate their privileges. ([CVE-2010-4649](#)²³, Important)
- * An integer signedness flaw in `drm_modeset_ctl()` could allow a local, unprivileged user to cause a denial of service or escalate their privileges. ([CVE-2011-1013](#)²⁴, Important)
- * The Radeon GPU drivers in the Linux kernel were missing sanity checks for the Anti Aliasing (AA) resolve register values which could allow a local, unprivileged user to cause a denial of service or escalate their privileges on systems using a graphics card from the ATI Radeon R300, R400, or R500 family of cards. ([CVE-2011-1016](#)²⁵, Important)
- * A flaw in `dccp_rcv_state_process()` could allow a remote attacker to cause a denial of service, even when the socket was already closed. ([CVE-2011-1093](#)²⁶, Important)
- * A flaw in the Linux kernel's Stream Control Transmission Protocol (SCTP) implementation could allow a remote attacker to cause a denial of service if the `sysctl net.sctp.addip_enable` and `auth_enable` variables were turned on (they are off by default). ([CVE-2011-1573](#)²⁷, Important)
- * A memory leak in the `inotify_init()` system call. In some cases, it could leak a group, which could allow a local, unprivileged user to eventually cause a denial of service. ([CVE-2010-4250](#)²⁸, Moderate)
- * A missing validation of a null-terminated string data structure element in `bnep_sock_ioctl()` could allow a local user to cause an information leak or a denial of service. ([CVE-2011-1079](#)²⁹, Moderate)

²² <https://rhn.redhat.com/errata/RHSA-2011-0498.html>

²³ <https://www.redhat.com/security/data/cve/CVE-2010-4649.html>

²⁴ <https://www.redhat.com/security/data/cve/CVE-2011-1013.html>

²⁵ <https://www.redhat.com/security/data/cve/CVE-2011-1016.html>

²⁶ <https://www.redhat.com/security/data/cve/CVE-2011-1093.html>

²⁷ <https://www.redhat.com/security/data/cve/CVE-2011-1573.html>

²⁸ <https://www.redhat.com/security/data/cve/CVE-2010-4250.html>

²⁹ <https://www.redhat.com/security/data/cve/CVE-2011-1079.html>

* An information leak in `bcm_connect()` in the Controller Area Network (CAN) Broadcast Manager implementation could allow a local, unprivileged user to leak kernel mode addresses in `/proc/net/can-bcm`. ([CVE-2010-4565](#)³⁰, Low)

* A flaw was found in the Linux kernel's Integrity Measurement Architecture (IMA) implementation. When SELinux was disabled, adding an IMA rule which was supposed to be processed by SELinux would cause `ima_match_rules()` to always succeed, ignoring any remaining rules. ([CVE-2011-0006](#)³¹, Low)

* A missing initialization flaw in the XFS file system implementation could lead to an information leak. ([CVE-2011-0711](#)³², Low)

* Buffer overflow flaws in `snd_usb_caiaq_audio_init()` and `snd_usb_caiaq_midi_init()` could allow a local, unprivileged user with access to a Native Instruments USB audio device to cause a denial of service or escalate their privileges. ([CVE-2011-0712](#)³³, Low)

* The `start_code` and `end_code` values in `/proc/<PID>/stat` were not protected. In certain scenarios, this flaw could be used to defeat Address Space Layout Randomization (ASLR). ([CVE-2011-0726](#)³⁴, Low)

* A flaw in `dev_load()` could allow a local user who has the `CAP_NET_ADMIN` capability to load arbitrary modules from `/lib/modules/`, instead of only `netdev` modules. ([CVE-2011-1019](#)³⁵, Low)

* A flaw in `ib_uverbs_poll_cq()` could allow a local, unprivileged user to cause an information leak. ([CVE-2011-1044](#)³⁶, Low)

* A missing validation of a null-terminated string data structure element in `do_replace()` could allow a local user who has the `CAP_NET_ADMIN` capability to cause an information leak. ([CVE-2011-1080](#)³⁷, Low)

Red Hat would like to thank Vegard Nossum for reporting CVE-2010-4250; Vasilij Kulikov for reporting CVE-2011-1079, CVE-2011-1019, and CVE-2011-1080; Dan Rosenberg for reporting CVE-2010-4565 and CVE-2011-0711; Rafael Dominguez Vega for reporting CVE-2011-0712; and Kees Cook for reporting CVE-2011-0726.

Bug fixes:

BZ#[659572](#)³⁸

A flaw was found in the Linux kernel where, if used in conjunction with another flaw that can result in a kernel Oops, could possibly lead to privilege escalation. It does not affect Red Hat Enterprise Linux 6 as the `sysctl panic_on_oops` variable is turned on by default. However, as a preventive measure if the variable is turned off by an administrator, this update addresses the issue. Red Hat would like to thank Nelson Elhage for reporting this vulnerability.

BZ#[694073](#)³⁹

Under some circumstances, faulty logic in the system BIOS could report that ASPM (Active State Power Management) was not supported on the system, but leave ASPM enabled on a device.

³⁰ <https://www.redhat.com/security/data/cve/CVE-2010-4565.html>

³¹ <https://www.redhat.com/security/data/cve/CVE-2011-0006.html>

³² <https://www.redhat.com/security/data/cve/CVE-2011-0711.html>

³³ <https://www.redhat.com/security/data/cve/CVE-2011-0712.html>

³⁴ <https://www.redhat.com/security/data/cve/CVE-2011-0726.html>

³⁵ <https://www.redhat.com/security/data/cve/CVE-2011-1019.html>

³⁶ <https://www.redhat.com/security/data/cve/CVE-2011-1044.html>

³⁷ <https://www.redhat.com/security/data/cve/CVE-2011-1080.html>

³⁸ https://bugzilla.redhat.com/show_bug.cgi?id=659572

³⁹ https://bugzilla.redhat.com/show_bug.cgi?id=694073

This could lead to AER (Advanced Error Reporting) errors that the kernel was unable to handle. With this update, the kernel proactively disables ASPM on devices when the BIOS reports that ASPM is not supported, safely eliminating the aforementioned issues.

BZ#696487⁴⁰

Prior to this update, adding a bond over a bridge inside a virtual guest caused the kernel to crash due to a NULL dereference. This update improves the tests for the presence of VLANs configured above bonding (additionally, this update fixes a regression introduced by the patch for BZ#633571⁴¹). The new logic determines whether a registration has occurred, instead of testing that the internal `vlan_list` of a bond is empty. Previously, the system panicked and crashed when `vlan_list` was not empty, but the `vlggrp` pointer was still NULL.

BZ#698109⁴²

During light or no network traffic, the active-backup interface bond using ARP monitoring with validation could go down and return due to an overflow or underflow of system timer interrupt ticks (jiffies). With this update, the jiffies calculation issues have been fixed and a bond interface works as expected.

BZ#691777⁴³

In certain network setups (specifically, using VLAN on certain NICs where packets are sent through the VLAN GRO rx path), sending packets from an active ethernet port to another inactive ethernet port could affect the network's bridge and cause the bridge to acquire a wrong bridge port. This resulted in all packets not being passed along in the network. With this update, the underlying source code has been modified to address this issue, and network traffic works as expected.

BZ#698114⁴⁴, BZ#696889⁴⁵

Deleting a SCSI (Small Computer System Interface) device attached to a device handler caused applications running in user space, which were performing I/O operations on that device, to become unresponsive. This was due to the fact that the SCSI device handler's activation did not propagate the SCSI device deletion via an error code and a callback to the Device-Mapper Multipath. With this update, deletion of a SCSI device attached to a device handler is properly handled and no longer causes certain applications to become unresponsive.

BZ#683440⁴⁶

Systems Management Applications using the `libsmbios` package could become unresponsive on Dell PowerEdge servers (specifically, Dell PowerEdge 2970 and Dell PowerEdge SC1435). The `dcdbas` driver can perform an I/O write operation which causes an SMI (System Management Interrupt) to occur. However, the SMI handler processed the SMI well after the `outb` function was processed, which caused random failures resulting in the aforementioned hang. With this update, the underlying source code has been modified to address this issue, and systems management applications using the `libsmbios` package no longer become unresponsive.

BZ#670850⁴⁷

Invoking an EFI (Extensible Firmware Interface) call caused a restart or a failure to boot to occur on a system with more than 512GB of memory because the EFI page tables did not map the

⁴⁰ https://bugzilla.redhat.com/show_bug.cgi?id=696487

⁴¹ https://bugzilla.redhat.com/show_bug.cgi?id=633571

⁴² https://bugzilla.redhat.com/show_bug.cgi?id=698109

⁴³ https://bugzilla.redhat.com/show_bug.cgi?id=691777

⁴⁴ https://bugzilla.redhat.com/show_bug.cgi?id=698114

⁴⁵ https://bugzilla.redhat.com/show_bug.cgi?id=696889

⁴⁶ https://bugzilla.redhat.com/show_bug.cgi?id=683440

⁴⁷ https://bugzilla.redhat.com/show_bug.cgi?id=670850

whole kernel space. EFI page tables used only one PGD (Page Global Directory) entry to map the kernel space; thus, virtual addresses higher than `PAGE_OFFSET + 512GB` could not be accessed. With this update, EFI page tables map the whole kernel space.

BZ#683820⁴⁸

Enabling the Header Splitting mode on all Intel 82599 10 Gigabit Ethernet hardware could lead to unpredictable behavior. With this update, the Header Splitting mode is never enabled on the aforementioned hardware.

BZ#670114⁴⁹

The `ixgbe` driver has been upgraded to upstream version 3.0.12, which provides a number of bug fixes and enhancements over the previous version.

BZ#670110⁵⁰

If an Intel 82598 10 Gigabit Ethernet Controller was configured in a way that caused peer-to-peer traffic to be sent to the Intel X58 I/O hub (IOH), a PCIe credit starvation problem occurred. As a result, the system would hang. With this update, the system continues to work and does not hang.

BZ#683817⁵¹

The ALSA HDA audio driver has been updated to improve support for new chipsets and HDA audio codecs.

BZ#689341⁵²

A buffer overflow flaw was found in the Linux kernel's Cluster IP hashmark target implementation. A local, unprivileged user could trigger this flaw and cause a local denial of service by editing files in the `/proc/net/ipt_CLUSTERIP/` directory. Note: On Red Hat Enterprise 6, only root can write to files in the `/proc/net/ipt_CLUSTERIP/` directory by default. This update corrects this issue as a preventative measure in case an administrator has changed the permissions on these files. Red Hat would like to thank Vasilij Kulikov for reporting this issue.

BZ#684275⁵³

Using the `pam_tty_audit.so` module (which enables or disables TTY auditing for specified users) in the `/etc/pam.d/sudo` file and in the `/etc/pam.d/system-auth` file when the audit package is not installed resulted in soft lock-ups on CPUs. As a result, the kernel became unresponsive. This was due to the kernel exiting immediately after TTY auditing was disabled, without emptying the buffer, which caused the kernel to spin in a loop, copying 0 bytes at each iteration and attempting to push each time without any effect. With this update, a locking mechanism is introduced to prevent the aforementioned behavior.

BZ#679306⁵⁴

Prior to this update, a collection of world-writable `sysfs` and `procfs` files allowed an unprivileged user to change various settings, change device hardware registers, and load certain firmware. With this update, permissions for these files have been changed.

⁴⁸ https://bugzilla.redhat.com/show_bug.cgi?id=683820

⁴⁹ https://bugzilla.redhat.com/show_bug.cgi?id=670114

⁵⁰ https://bugzilla.redhat.com/show_bug.cgi?id=670110

⁵¹ https://bugzilla.redhat.com/show_bug.cgi?id=683817

⁵² https://bugzilla.redhat.com/show_bug.cgi?id=689341

⁵³ https://bugzilla.redhat.com/show_bug.cgi?id=684275

⁵⁴ https://bugzilla.redhat.com/show_bug.cgi?id=679306

BZ#694186⁵⁵

A previously introduced patch could cause kswapd (the kernel's memory reclaim daemon) to enter an infinite loop, consuming 100% of the CPU it is running on. This happened because kswapd incorrectly stayed awake for an unreclaimable zone. This update addresses this issue, and kswapd no longer consumes 100% of the CPU it is running on.

BZ#695322⁵⁶

If an error occurred during an I/O operation, the SCSI driver reset the megaraid_sas controller to restore it to normal state. However, on Red Hat Enterprise Linux 6, the waiting time to allow a full reset completion for the megaraid_sas controller was too short. The driver incorrectly recognized the controller as stalled, and, as a result, the system stalled as well. With this update, more time is given to the controller to properly restart, thus, the controller operates as expected after being reset.

Enhancement:

BZ#683810⁵⁷

This update provides VLAN null tagging support (**VLAN ID 0** can be used in tags).

Users should upgrade to these updated packages, which contain backported patches to correct these issues and add this enhancement. The system must be rebooted for this update to take effect.

B.1.3. RHSA-2011:0421 - Important: kernel security and bug fix update



Important

This update has already been released as the security errata [RHSA-2011:0421](https://rhn.redhat.com/errata/RHSA-2011-0421.html)⁵⁸

Updated kernel packages that resolve several security issues and fix various bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links after each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security fixes:

* A flaw was found in the `sctp_icmp_proto_unreachable()` function in the Linux kernel's Stream Control Transmission Protocol (SCTP) implementation. A remote attacker could use this flaw to cause a denial of service. ([CVE-2010-4526](https://www.redhat.com/security/data/cve/CVE-2010-4526.html)⁵⁹, Important)

* A missing boundary check was found in the `dvb_ca_ioctl()` function in the Linux kernel's `av7110` module. On systems that use old DVB cards that require the `av7110` module, a local, unprivileged

⁵⁵ https://bugzilla.redhat.com/show_bug.cgi?id=694186

⁵⁶ https://bugzilla.redhat.com/show_bug.cgi?id=695322

⁵⁷ https://bugzilla.redhat.com/show_bug.cgi?id=683810

⁵⁸ <https://rhn.redhat.com/errata/RHSA-2011-0421.html>

⁵⁹ <https://www.redhat.com/security/data/cve/CVE-2010-4526.html>

Technical Notes

user could use this flaw to cause a denial of service or escalate their privileges. ([CVE-2011-0521](https://www.redhat.com/security/data/cve/CVE-2011-0521)⁶⁰, Important)

* A race condition was found in the way the Linux kernel's InfiniBand implementation set up new connections. This could allow a remote user to cause a denial of service. ([CVE-2011-0695](https://www.redhat.com/security/data/cve/CVE-2011-0695)⁶¹, Important)

* A heap overflow flaw in the `iowarrior_write()` function could allow a user with access to an IO-Warrior USB device, that supports more than 8 bytes per report, to cause a denial of service or escalate their privileges. ([CVE-2010-4656](https://www.redhat.com/security/data/cve/CVE-2010-4656)⁶², Moderate)

* A flaw was found in the way the Linux Ethernet bridge implementation handled certain IGMP (Internet Group Management Protocol) packets. A local, unprivileged user on a system that has a network interface in an Ethernet bridge could use this flaw to crash that system. ([CVE-2011-0716](https://www.redhat.com/security/data/cve/CVE-2011-0716)⁶³, Moderate)

* A NULL pointer dereference flaw was found in the Generic Receive Offload (GRO) functionality in the Linux kernel's networking implementation. If both GRO and promiscuous mode were enabled on an interface in a virtual LAN (VLAN), it could result in a denial of service when a malformed VLAN frame is received on that interface. ([CVE-2011-1478](https://www.redhat.com/security/data/cve/CVE-2011-1478)⁶⁴, Moderate)

* A missing initialization flaw in the Linux kernel could lead to an information leak. ([CVE-2010-3296](https://www.redhat.com/security/data/cve/CVE-2010-3296)⁶⁵, Low)

* A missing security check in the Linux kernel's implementation of the `install_special_mapping` routine could allow a local, unprivileged user to bypass the `mmap_min_addr` protection mechanism. ([CVE-2010-4346](https://www.redhat.com/security/data/cve/CVE-2010-4346)⁶⁶, Low)

* A logic error in the `orinoco_ioctl_set_auth()` function in the Linux kernel's ORiNOCO wireless extensions support implementation could render TKIP countermeasures ineffective when it is enabled, as it enabled the card instead of shutting it down. ([CVE-2010-4648](https://www.redhat.com/security/data/cve/CVE-2010-4648)⁶⁷, Low)

* A missing initialization flaw was found in the `ethtool_get_regs()` function in the Linux kernel's ethtool IOCTL handler. A local user who has the `CAP_NET_ADMIN` capability could use this flaw to cause an information leak. ([CVE-2010-4655](https://www.redhat.com/security/data/cve/CVE-2010-4655)⁶⁸, Low)

* An information leak was found in the Linux kernel's `task_show_regs()` implementation. On IBM S/390 systems, a local, unprivileged user could use this flaw to read `/proc/<PID>/status` files, allowing them to discover the CPU register values of processes. ([CVE-2011-0710](https://www.redhat.com/security/data/cve/CVE-2011-0710)⁶⁹, Low)

Red Hat would like to thank Jens Kuehnel for reporting CVE-2011-0695; Kees Cook for reporting CVE-2010-4656 and CVE-2010-4655; Dan Rosenberg for reporting CVE-2010-3296; and Tavis Ormandy for reporting CVE-2010-4346.

Bug fixes:

⁶⁰ <https://www.redhat.com/security/data/cve/CVE-2011-0521.html>

⁶¹ <https://www.redhat.com/security/data/cve/CVE-2011-0695.html>

⁶² <https://www.redhat.com/security/data/cve/CVE-2010-4656.html>

⁶³ <https://www.redhat.com/security/data/cve/CVE-2011-0716.html>

⁶⁴ <https://www.redhat.com/security/data/cve/CVE-2011-1478.html>

⁶⁵ <https://www.redhat.com/security/data/cve/CVE-2010-3296.html>

⁶⁶ <https://www.redhat.com/security/data/cve/CVE-2010-4346.html>

⁶⁷ <https://www.redhat.com/security/data/cve/CVE-2010-4648.html>

⁶⁸ <https://www.redhat.com/security/data/cve/CVE-2010-4655.html>

⁶⁹ <https://www.redhat.com/security/data/cve/CVE-2011-0710.html>

BZ#678484⁷⁰

The `bnx2i` driver could cause a system crash on IBM POWER7 systems. The driver's page tables were not set up properly on Big Endian machines, causing extended error handling (EEH) errors on PowerPC machines. With this update, the page tables are properly set up and a system crash no longer occurs in the aforementioned case.

BZ#678485⁷¹

On platforms using an Intel 7500 or an Intel 5500 chipset (or their derivatives), occasionally, a VT-d specification defined error occurred in the `kdump` kernel (the second kernel). As a result of the VT-d error, on some platforms, an SMI (System Management Interrupt) was issued and the system became unresponsive. With this update, a VT-d error is properly handled so that an SMI is no longer issued, and the system no longer hangs.

BZ#678558⁷²

Using a `virtio` serial port from an application, filling it until the `write` command returns `-EAGAIN` and then executing a `select` command for the `write` command, caused the `select` command to not return any values when using the `virtio` serial port in a *non-blocking* mode. When used in *blocking* mode, the `write` command waited until the host indicated it had used up the buffers. This was due to the fact that the poll operation waited for the `port->waitqueue` pointer; however, nothing woke the `waitqueue` when there was room again in the queue. With this update, the queue is woken via host notifications so that buffers consumed by the host can be reclaimed, the queue freed, and the application `write` operations may proceed again.

BZ#678559⁷³

Prior to this update, user space could submit (using the `write()` operation) a buffer with zero length to be written to the host, causing the qemu hypervisor instance running on that host to crash. This was caused by the `write()` operation triggering a `virtqueue` event on the host, causing a NULL buffer to be accessed. With this update, user space is no longer allowed to submit zero-sized buffers and the aforementioned crash no longer occur.

BZ#678561⁷⁴

Applications and agents using `virtio` serial ports would block messages even though there were messages queued up and ready to be read in the `virtqueue`. This was due to `virtio_console`'s poll function checking whether a port was NULL to determine if a read operation would result in a block of the port. However, in some cases, a port can be NULL even though there are buffers left in the `virtqueue` to be read. This update introduces a more sophisticated method of checking whether a port contains any data; thus, preventing queued up messages from being incorrectly blocked.

BZ#678562⁷⁵

If a host was slow in reading data or did not read data at all, blocking `write()` calls not only blocked the program that called the `write()` call but also the entire guest. This was caused by the `write()` calls waiting until an acknowledgment that the data consumed was received from the host. With this update, `write()` calls no longer wait for such acknowledgment: control is immediately returned to the user space application. This ensures that even if the host is busy processing other data or is not consuming data at all, the guest is not blocked.

⁷⁰ https://bugzilla.redhat.com/show_bug.cgi?id=678484

⁷¹ https://bugzilla.redhat.com/show_bug.cgi?id=678485

⁷² https://bugzilla.redhat.com/show_bug.cgi?id=678558

⁷³ https://bugzilla.redhat.com/show_bug.cgi?id=678559

⁷⁴ https://bugzilla.redhat.com/show_bug.cgi?id=678561

⁷⁵ https://bugzilla.redhat.com/show_bug.cgi?id=678562

BZ#678996⁷⁶

An implementation of the SHA (Secure Hash Algorithm) hashing algorithm for the IBM System z architecture did not produce correct hashes and could potentially cause memory corruption due to broken partial block handling. A partial block could break when it was followed by an update which filled it with leftover bytes. Instead of storing the new leftover bytes at the start of the buffer, they were stored immediately after the previous partial block. With this update, the index pointer is reset, thus resolving the aforementioned partial block handling issue.

BZ#680080⁷⁷

Prior to this update, performing live migration back and forth during guest installation with network adapters based on the 8168c chipset or the 8111c chipset triggered an `rtl8169_interrupt` hang due to a RxFIFO overflow. With this update, infinite loops in the IRQ (Interrupt Request) handler caused by RxFIFO overflows are prevented and the aforementioned hang no longer occurs.

BZ#683442⁷⁸

Reading the `/proc/vmcore` file was previously significantly slower on a Red Hat Enterprise Linux 6 system when compared to a Red Hat Enterprise Linux 5 system. This update enables caching of memory accesses; reading of the `/proc/vmcore` file is now noticeably faster.

BZ#683445⁷⁹

Reading the `/proc/vmcore` file on a Red Hat Enterprise Linux 6 system was not optimal because it did not always take advantage of reading through the cached memory. With this update, access to the `/dev/oldmem` device in the `/proc/vmcore` file is cached, resulting in faster copying to user space.

BZ#683781⁸⁰

Migrating a guest could have resulted in dirty values for the guest being retained in memory, which could have caused both the guest and qemu to crash. The trigger for this was memory pages being both write-protected and dirty simultaneously. With this update, memory pages in the current bitmap are either dirty or write-protected when migrating a guest, with the result that neither qemu nor guest operating systems crash following a migration.

BZ#683783⁸¹

While not mandated by any specification, Linux systems rely on NMIs (**N**on-**m**askable **I**nterrupts) being blocked by an IF-enabling (**I**nterrupt **F**lag) STI instruction (an x86 instruction that enables interrupts; **S**et **I**nterrupts); this is also the common behavior of all known hardware. Prior to this update, kernel panic could occur on guests using NMIs extensively (for example, a Linux system with the `nmi_watchdog` kernel parameter enabled). With this update, an NMI is disallowed when interrupts are blocked by an STI. This is done by checking for the condition and requesting an interrupt window exit if it occurs. As a result, kernel panic no longer occurs.

BZ#683812⁸²

Under certain circumstances, a kernel thread that handles incoming messages from a server could unexpectedly exit by itself. As a result, the kernel thread would free some data structures which could then be referenced by another data structure, resulting in a kernel panic. With this

⁷⁶ https://bugzilla.redhat.com/show_bug.cgi?id=678996

⁷⁷ https://bugzilla.redhat.com/show_bug.cgi?id=680080

⁷⁸ https://bugzilla.redhat.com/show_bug.cgi?id=683442

⁷⁹ https://bugzilla.redhat.com/show_bug.cgi?id=683445

⁸⁰ https://bugzilla.redhat.com/show_bug.cgi?id=683781

⁸¹ https://bugzilla.redhat.com/show_bug.cgi?id=683783

⁸² https://bugzilla.redhat.com/show_bug.cgi?id=683812

update, kernel threads no longer unexpectedly exit; thus, kernel panic no longer occurs in the aforementioned case.

BZ#683814⁸³

Operating in the FIP (FCoE Initialization Protocol) mode and performing operations that bring up ports could cause the `fcoe.ko` and `fnic.ko` modules to not be able to re-login when a port was brought back up. This was due to a bug in the FCoE (Fiber Channel over Ethernet) layer causing improper handling of FCoE LOGO frames while in the FIP mode. With this update, FCoE LOGO frames are properly handled when in the FIP mode and the `fcoe.ko` and `fnic.ko` modules no longer fail to re-login.

BZ#683815⁸⁴

If a CPU is set offline, the `nohz_load_balancer` CPU is updated. However, under certain circumstances, the `nohz_load_balancer` CPU would not be updated, causing the offlined CPU to be enqueued with various timers which never expired. As a result, the system could become unresponsive. With this update, the `nohz_load_balancer` CPU is always updated; systems no longer become unresponsive.

BZ#683822⁸⁵

The kernel syslog contains debugging information that is often useful during exploitation of other vulnerabilities such as kernel heap addresses. With this update, a new **CONFIG_SECURITY_DMESG_RESTRICT** option has been added to `config-generic-rhel` which prevents unprivileged users from reading the kernel syslog. This option is by default turned off (0), which means no restrictions.

BZ#684129⁸⁶

Prior to this update, the default VF (Virtual Function) configuration was not restrictive enough. With this update, VFs only accept broadcast and multicast frames and do not accept frames from the unicast MAC address table. Restrictions are now also properly set on what can be received when the device is put in promiscuous mode. A hardware limitation was also discovered that prevented the system from properly receiving certain FCoE (Fibre Channel over Ethernet) protocol frames of a specific size. A buffer management change now allows these frames to be properly received.

BZ#684266⁸⁷

PowerPC systems having more than 1 TB of RAM could randomly crash or become unresponsive due to an incorrect setup of the Segment Lookaside Buffer (SLB) entry for the kernel stack. With this update, the SLB entry is properly set up.

BZ#684267⁸⁸

On IBM System z systems, user space programs could access the `/dev/mem` file (which contains an image of main memory), where an accidental memory (write) access could potentially be harmful. To restrict access to memory from user space through the `/dev/mem` file, the **CONFIG_STRICT_DEVMEM** configuration option has been enabled for the default kernel. The `kdump` and `debug` kernels have this option switched off by default.

⁸³ https://bugzilla.redhat.com/show_bug.cgi?id=683814

⁸⁴ https://bugzilla.redhat.com/show_bug.cgi?id=683815

⁸⁵ https://bugzilla.redhat.com/show_bug.cgi?id=683822

⁸⁶ https://bugzilla.redhat.com/show_bug.cgi?id=684129

⁸⁷ https://bugzilla.redhat.com/show_bug.cgi?id=684266

⁸⁸ https://bugzilla.redhat.com/show_bug.cgi?id=684267

BZ#684268⁸⁹

Intensive usage of resources on a guest lead to a failure of networking on that guest: packets could no longer be received. The failure occurred when a DMA (Direct Memory Access) ring was consumed before NAPI (New API; an interface for networking devices which makes use of interrupt mitigation techniques) was enabled which resulted in a failure to receive the next interrupt request. The regular interrupt handler was not affected in this situation (because it can process packets in-place), however, the OOM (Out Of Memory) handler did not detect the aforementioned situation and caused networking to fail. With this update, NAPI is subsequently scheduled for each `napi_enable` operation; thus, networking no longer fails under the aforementioned circumstances.

Users should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

B.1.4. RHSA-2011:0283 - Moderate: kernel security, bug fix and enhancement update



Important

This update has already been released as the security errata [RHSA-2011:0283](#)⁹⁰

Updated kernel packages that resolve several security issues, fix various bugs and add enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links after each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security fixes:

* A divide-by-zero flaw was found in the `tcp_select_initial_window()` function in the Linux kernel's TCP/IP protocol suite implementation. A local, unprivileged user could use this flaw to trigger a denial of service by calling `setsockopt()` with certain options. ([CVE-2010-4165](#)⁹¹, Moderate)

* A use-after-free flaw in the `mprotect()` system call in the Linux kernel could allow a local, unprivileged user to cause a local denial of service. ([CVE-2010-4169](#)⁹², Moderate)

* A flaw was found in the Linux kernel `execve()` system call implementation. A local, unprivileged user could cause large amounts of memory to be allocated but not visible to the OOM (Out of Memory) killer, triggering a denial of service. ([CVE-2010-4243](#)⁹³, Moderate)

Red Hat would like to thank Steve Chen for reporting CVE-2010-4165, and Brad Spengler for reporting CVE-2010-4243.

⁸⁹ https://bugzilla.redhat.com/show_bug.cgi?id=684268

⁹⁰ <https://rhn.redhat.com/errata/RHSA-2011-0283.html>

⁹¹ <https://www.redhat.com/security/data/cve/CVE-2010-4165.html>

⁹² <https://www.redhat.com/security/data/cve/CVE-2010-4169.html>

⁹³ <https://www.redhat.com/security/data/cve/CVE-2010-4243.html>

Bug fixes:

BZ#[652720](https://bugzilla.redhat.com/show_bug.cgi?id=652720)⁹⁴

Prior to this update, a guest could use the `poll()` function to find out whether the host-side connection was open or closed. However, with a `SIGIO` signal, this can be done asynchronously, without having to explicitly poll each port. With this update, a `SIGIO` signal is sent for any host connect/disconnect events. Once the `SIGIO` signal is received, the open/close status of `virtio-serial` ports can be obtained using the `poll()` system call.

BZ#[658854](https://bugzilla.redhat.com/show_bug.cgi?id=658854)⁹⁵

A Red Hat Enterprise Linux 6.0 host (with root on a local disk) with `dm-multipath` configured on multiple LUNs (Logical Unit Number) hit kernel panic (at `scsi_error_handler`) with target controller faults during an I/O operation on the `dm-multipath` devices. This was caused by `multipath` using the `blk_abort_queue()` function to allow lower latency path deactivation. The call to `blk_abort_queue` proved to be unsafe due to a race (between `blk_abort_queue` and `scsi_request_fn`). With this update, the race has been resolved and kernel panic no longer occurs on Red Hat Enterprise Linux 6.0 hosts.

BZ#[658891](https://bugzilla.redhat.com/show_bug.cgi?id=658891)⁹⁶

Prior to this update, running context-switch intensive workloads on KVM guests resulted in a large number of exits (`kvm_exit`) due to control register (CR) accesses by the guest, thus, resulting in poor performance. This update includes a number of optimizations which allow the guest not to exit to the hypervisor in the aforementioned case and improve the overall performance.

BZ#[659610](https://bugzilla.redhat.com/show_bug.cgi?id=659610)⁹⁷

Handling ALUA (Asymmetric Logical Unit Access) *transitioning* states did not work properly due to a faulty SCSI (Small Computer System Interface) ALUA handler. With this update, optimized state transitioning prevents the aforementioned behavior.

BZ#[660590](https://bugzilla.redhat.com/show_bug.cgi?id=660590)⁹⁸

Prior to this update, when using Red Hat Enterprise Linux 6 with a `qla4xxx` driver and FC (Fibre Channel) drivers using the `fc` class, a device might have been put in the *offline* state due to a transport problem. Once the transport problem was resolved, the device was not usable until a user manually corrected the state. This update enables the transition from the *offline* state to the *running* state, thus, fixing the problem.

BZ#[661667](https://bugzilla.redhat.com/show_bug.cgi?id=661667)⁹⁹

The `zfcpdump` tool was not able to mount `ext4` file systems. Because `ext4` is the default file system on Red Hat Enterprise Linux 6, with this update, `ext4` file system support was added for the `zfcpdump` tool.

BZ#[661725](https://bugzilla.redhat.com/show_bug.cgi?id=661725)¹⁰⁰

The `zfcpdump` tool was not able to mount `ext2` file systems. With this update, `ext2` file system support was added for the `zfcpdump` tool.

⁹⁴ https://bugzilla.redhat.com/show_bug.cgi?id=652720

⁹⁵ https://bugzilla.redhat.com/show_bug.cgi?id=658854

⁹⁶ https://bugzilla.redhat.com/show_bug.cgi?id=658891

⁹⁷ https://bugzilla.redhat.com/show_bug.cgi?id=659610

⁹⁸ https://bugzilla.redhat.com/show_bug.cgi?id=660590

⁹⁹ https://bugzilla.redhat.com/show_bug.cgi?id=661667

¹⁰⁰ https://bugzilla.redhat.com/show_bug.cgi?id=661725

BZ#661730¹⁰¹

The *lock reclaim* operation on a Red Hat Enterprise Linux 6 NFSv4 client did not work properly when, after a server reboot, an I/O operation which resulted in a *STALE_STATEID* response was performed before the *RENEW* call was sent to the server. This behavior was caused due to the improper use of the state flags. While investigating this bug, a different bug was discovered in the *state recovery* operation which resulted in a reclaim thread looping in the `nfs4_reclaim_open_state()` function. With this update, both operations have been fixed and work as expected.

BZ#661731¹⁰²

Prior to this update, the **execve** utility exhibited the following flaw. When an argument and any environment data were copied from an old task's user stack to the user stack of a newly-execve'd task, the kernel would not allow the process to be interrupted or rescheduled. Therefore, when the argument or environment string data was (abnormally) large, there was no "interactivity" with the process while the `execve()` function was transferring the data. With this update, fatal signals (like **CTRL+c**) can now be received and handled and a process is allowed to yield to higher priority processes during the data transfer.

BZ#661732¹⁰³

The *memory* cgroup controller has its own Out of Memory routine (OOM killer) and kills a process at an OOM event. However, a race condition could cause the `pagefault_out_of_memory` function to be called after the *memory* cgroup's OOM. This invoked the generic OOM killer and a *panic_on_oom* could occur. With this update, only the *memory* cgroup's OOM killer is invoked and used to kill a process should an OOM occur.

BZ#661737¹⁰⁴

In some cases, under a small system load involve some I/O operation, processes started to lock up in the D state (that is, became unresponsive). The system load could in some cases climb steadily. This was due to the way the event channel IRQ (Interrupt Request) was set up. Xen events behave like edge-triggered IRQs, however, the kernel was setting them up as level-triggered IRQs. As a result, any action using Xen event channels could lock up a process in the D state. With this update, the handling has been changed from edge-triggered IRQs to level-triggered IRQs and process no longer lock up in the D state.

BZ#662049¹⁰⁵

When an **scsi** command timed out and the `fcoe/libfc` driver aborted the command, a race could occur during the clean-up of the command which could result in kernel panic. With this update, the locking mechanism in the clean-up and abort paths was modified, thus, fixing the aforementioned issue.

BZ#662050¹⁰⁶

The lack of synchronization between the clearing of the `QUEUE_FLAG_CLUSTER` flag and the setting of the `no_cluster` flag in the `queue_limits` variable caused corruption of data. Note that this issue only occurred on hardware that did not support segment merging (that is, clustering). With this update, the synchronization between the aforementioned flags works as expected, thus, corruption of data no longer occurs.

¹⁰¹ https://bugzilla.redhat.com/show_bug.cgi?id=661730

¹⁰² https://bugzilla.redhat.com/show_bug.cgi?id=661731

¹⁰³ https://bugzilla.redhat.com/show_bug.cgi?id=661732

¹⁰⁴ https://bugzilla.redhat.com/show_bug.cgi?id=661737

¹⁰⁵ https://bugzilla.redhat.com/show_bug.cgi?id=662049

¹⁰⁶ https://bugzilla.redhat.com/show_bug.cgi?id=662050

BZ#662721¹⁰⁷

The `virtio-console` device did not handle the `hot-unplug` operation properly. As a result, `virtio-console` could access the memory outside the driver's memory area and cause kernel panic on the guest. With this update, multiple fixes to the `virtio-console` device resolved this issue and the `hot-unplug` operation works as expected.

BZ#662921¹⁰⁸

Prior to this update, running the `hwclock --systemd` command could halt a running system. This was due to the interrupt transactions being looped back from a local IOH (Input/Output Hub), through the IOH to a local CPU (erroneously), which caused a conflict with I/O port operations and other transactions. With this update, the conflicts are avoided and the system continues to run after executing the `hwclock --systemd` command.

BZ#666797¹⁰⁹

An I/O operation could fast fail when using Device-Mapper Multipathing (`dm-multipath`) if the I/O operation could be retried by the `scsi` layer. This prevented the `multipath` layer from starting its error recovery procedure and resulted in unnecessary log messages in the appropriate log files. This update includes a number of optimizations that resolve the aforementioned issue.

BZ#670421¹¹⁰

Outgoing packets were not fragmented after receiving the `icmpv6 pkt-too-big` message when using the IPsecv6 tunnel mode. This was due to the lack of IPv6 fragmentation support over an IPsec tunnel. With this update, IPv6 fragmentation is fully supported and works as expected when using the IPsecv6 tunnel mode.

BZ#671342¹¹¹

Bonding, when operating in the ARP monitoring mode, made erroneous assumptions regarding the ownership of ARP frames when it received them for processing. Specifically, it was assumed that the bonding driver code was the only execution context which had access to the ARP frames network buffer data. As a result, an operation was attempted on the said buffer (specifically, to modify the size of the data buffer) which was forbidden by the kernel when a buffer was shared among several execution contexts. The result of such an operation on a shared buffer could lead to data corruption. Consequently, trying to prevent the corruption, the kernel panicked. This `shared state` in the network buffer could be forced to occur, for example, when running the `tcpdump` utility to monitor traffic on the bonding interface. Every buffer the bond interface received would be shared between the driver and the `tcpdump` process, thus, resulting in the aforementioned kernel panic. With this update, for the particular affected path in the bonding driver, each inbound frame is checked whether it is in the `shared state`. In case a buffer is shared, a private copy is made for exclusive use by the bonding driver, thus, preventing the kernel panic.

BZ#673978¹¹²

For a device that used a Target Portal Group (TPG) ID which occupied the full 2 bytes in the RTPG (Report Target Port Groups) response (with either byte exceeding the maximum value that may be stored in a signed char), the kernel's calculated TPG ID would never match the `group_id` that it should. As a result, this signed char overflow also caused the ALUA handler to incorrectly identify the Asymmetric Access State (AAS) of the specified device as well as incorrectly interpret

¹⁰⁷ https://bugzilla.redhat.com/show_bug.cgi?id=662721

¹⁰⁸ https://bugzilla.redhat.com/show_bug.cgi?id=662921

¹⁰⁹ https://bugzilla.redhat.com/show_bug.cgi?id=666797

¹¹⁰ https://bugzilla.redhat.com/show_bug.cgi?id=670421

¹¹¹ https://bugzilla.redhat.com/show_bug.cgi?id=671342

¹¹² https://bugzilla.redhat.com/show_bug.cgi?id=673978

the supported AAS of the target. With this update, the aforementioned issue has been addressed and no longer occurs.

Enhancements:

BZ#[674002](#)¹¹³

The `ixgbe` driver has been updated to address various FCoE (Fibre Channel over Ethernet) issues related to Direct Data Placement (FCoE DDP).

BZ#[664398](#)¹¹⁴

The `qla2xxx` driver for QLogic Fibre Channel Host Bus Adapters (HBAs) has been updated to upstream version 8.03.05.01.06.1-k0, which provides a number of bug fixes and enhancements over the previous version.

Users should upgrade to these updated packages, which contain backported patches to correct these issues, fix these bugs, and add these enhancements. The system must be rebooted for this update to take effect.

B.1.5. RHSA-2011:0007 - Important: kernel security and bug fix update



Important

This update has already been released as the security errata [RHSA-2011:0007](#)¹¹⁵

Updated kernel packages that fix multiple security issues and several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links after each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security fixes:

* Buffer overflow in `eCryptfs`. When `/dev/ecryptfs` has world writable permissions (which it does not, by default, on Red Hat Enterprise Linux 6), a local, unprivileged user could use this flaw to cause a denial of service or possibly escalate their privileges. ([CVE-2010-2492](#)¹¹⁶, Important)

* Integer overflow in the RDS protocol implementation could allow a local, unprivileged user to cause a denial of service or escalate their privileges. ([CVE-2010-3865](#)¹¹⁷, Important)

* Missing boundary checks in the PPP over L2TP sockets implementation could allow a local, unprivileged user to cause a denial of service or escalate their privileges. ([CVE-2010-4160](#)¹¹⁸, Important)

¹¹³ https://bugzilla.redhat.com/show_bug.cgi?id=674002

¹¹⁴ https://bugzilla.redhat.com/show_bug.cgi?id=664398

¹¹⁵ <https://rhn.redhat.com/errata/RHSA-2011-0007.html>

¹¹⁶ <https://www.redhat.com/security/data/cve/CVE-2010-2492.html>

¹¹⁷ <https://www.redhat.com/security/data/cve/CVE-2010-3865.html>

¹¹⁸ <https://www.redhat.com/security/data/cve/CVE-2010-4160.html>

- * NULL pointer dereference in the `igb` driver. If both Single Root I/O Virtualization (SR-IOV) and promiscuous mode were enabled on an interface using `igb`, it could result in a denial of service when a tagged VLAN packet is received on that interface. ([CVE-2010-4263](https://www.redhat.com/security/data/cve/CVE-2010-4263)¹¹⁹, Important)
- * Missing initialization flaw in the XFS file system implementation, and in the network traffic policing implementation, could allow a local, unprivileged user to cause an information leak. ([CVE-2010-3078](https://www.redhat.com/security/data/cve/CVE-2010-3078)¹²⁰, CVE-2010-3477, Moderate)
- * NULL pointer dereference in the Open Sound System compatible sequencer driver could allow a local, unprivileged user with access to `/dev/sequencer` to cause a denial of service. `/dev/sequencer` is only accessible to root and users in the audio group by default. ([CVE-2010-3080](https://www.redhat.com/security/data/cve/CVE-2010-3080)¹²¹, Moderate)
- * Flaw in the `ethtool` IOCTL handler could allow a local user to cause an information leak. ([CVE-2010-3861](https://www.redhat.com/security/data/cve/CVE-2010-3861)¹²², Moderate)
- * Flaw in `bcm_connect ()` in the Controller Area Network (CAN) Broadcast Manager. On 64-bit systems, writing the socket address may overflow the `procname` character array. ([CVE-2010-3874](https://www.redhat.com/security/data/cve/CVE-2010-3874)¹²³, Moderate)
- * Flaw in the module for monitoring the sockets of INET transport protocols could allow a local, unprivileged user to cause a denial of service. ([CVE-2010-3880](https://www.redhat.com/security/data/cve/CVE-2010-3880)¹²⁴, Moderate)
- * Missing boundary checks in the block layer implementation could allow a local, unprivileged user to cause a denial of service. ([CVE-2010-4162](https://www.redhat.com/security/data/cve/CVE-2010-4162)¹²⁵, CVE-2010-4163, CVE-2010-4668, Moderate)
- * NULL pointer dereference in the Bluetooth HCI UART driver could allow a local, unprivileged user to cause a denial of service. ([CVE-2010-4242](https://www.redhat.com/security/data/cve/CVE-2010-4242)¹²⁶, Moderate)
- * Flaw in the Linux kernel CPU time clocks implementation for the POSIX clock interface could allow a local, unprivileged user to cause a denial of service. ([CVE-2010-4248](https://www.redhat.com/security/data/cve/CVE-2010-4248)¹²⁷, Moderate)
- * Flaw in the garbage collector for AF_UNIX sockets could allow a local, unprivileged user to trigger a denial of service. ([CVE-2010-4249](https://www.redhat.com/security/data/cve/CVE-2010-4249)¹²⁸, Moderate)
- * Missing upper bound integer check in the AIO implementation could allow a local, unprivileged user to cause an information leak. ([CVE-2010-3067](https://www.redhat.com/security/data/cve/CVE-2010-3067)¹²⁹, Low)
- * Missing initialization flaws could lead to information leaks. ([CVE-2010-3298](https://www.redhat.com/security/data/cve/CVE-2010-3298)¹³⁰, [CVE-2010-3876](https://www.redhat.com/security/data/cve/CVE-2010-3876)¹³¹, [CVE-2010-4072](https://www.redhat.com/security/data/cve/CVE-2010-4072)¹³², [CVE-2010-4073](https://www.redhat.com/security/data/cve/CVE-2010-4073)¹³³, [CVE-2010-4074](https://www.redhat.com/security/data/cve/CVE-2010-4074)¹³⁴, [CVE-2010-4075](https://www.redhat.com/security/data/cve/CVE-2010-4075)¹³⁵, [CVE-2010-4077](https://www.redhat.com/security/data/cve/CVE-2010-4077)¹³⁶,

¹¹⁹ <https://www.redhat.com/security/data/cve/CVE-2010-4263.html>

¹²⁰ <https://www.redhat.com/security/data/cve/CVE-2010-3078.html>

¹²¹ <https://www.redhat.com/security/data/cve/CVE-2010-3080.html>

¹²² <https://www.redhat.com/security/data/cve/CVE-2010-3861.html>

¹²³ <https://www.redhat.com/security/data/cve/CVE-2010-3874.html>

¹²⁴ <https://www.redhat.com/security/data/cve/CVE-2010-3880.html>

¹²⁵ <https://www.redhat.com/security/data/cve/CVE-2010-4162.html>

¹²⁶ <https://www.redhat.com/security/data/cve/CVE-2010-4242.html>

¹²⁷ <https://www.redhat.com/security/data/cve/CVE-2010-4248.html>

¹²⁸ <https://www.redhat.com/security/data/cve/CVE-2010-4249.html>

¹²⁹ <https://www.redhat.com/security/data/cve/CVE-2010-3067.html>

¹³⁰ <https://www.redhat.com/security/data/cve/CVE-2010-3298.html>

¹³¹ <https://www.redhat.com/security/data/cve/CVE-2010-3876.html>

¹³² <https://www.redhat.com/security/data/cve/CVE-2010-4072.html>

¹³³ <https://www.redhat.com/security/data/cve/CVE-2010-4073.html>

¹³⁴ <https://www.redhat.com/security/data/cve/CVE-2010-4074.html>

¹³⁵ <https://www.redhat.com/security/data/cve/CVE-2010-4075.html>

[CVE-2010-4079](#)¹³⁷, [CVE-2010-4080](#)¹³⁸, [CVE-2010-4081](#)¹³⁹, [CVE-2010-4082](#)¹⁴⁰, [CVE-2010-4083](#)¹⁴¹, [CVE-2010-4158](#)¹⁴², Low)

* Missing initialization flaw in KVM could allow a privileged host user with access to `/dev/kvm` to cause an information leak. ([CVE-2010-4525](#)¹⁴³, Low)

Red Hat would like to thank Andre Osterhues for reporting CVE-2010-2492; Thomas Pollet for reporting CVE-2010-3865; Dan Rosenberg for reporting CVE-2010-4160, CVE-2010-3078, CVE-2010-3874, CVE-2010-4162, CVE-2010-4163, CVE-2010-3298, CVE-2010-4073, CVE-2010-4074, CVE-2010-4075, CVE-2010-4077, CVE-2010-4079, CVE-2010-4080, CVE-2010-4081, CVE-2010-4082, CVE-2010-4083, and CVE-2010-4158; Kosuke Tatsukawa for reporting CVE-2010-4263; Tavis Ormandy for reporting CVE-2010-3080 and CVE-2010-3067; Kees Cook for reporting CVE-2010-3861 and CVE-2010-4072; Nelson Elhage for reporting CVE-2010-3880; Alan Cox for reporting CVE-2010-4242; Vegard Nossum for reporting CVE-2010-4249; Vasily Kulikov for reporting CVE-2010-3876; and Stephan Mueller of atsec information security for reporting CVE-2010-4525.

Bug fixes:

BZ#[655122](#)¹⁴⁴

When building kernel modules against the full Red Hat Enterprise Linux 6 source tree (instead of just *kernel-devel*), modules would be signed by a locally generated key. However, Red Hat Enterprise Linux 6 refused to load modules created in this way as it did not recognize the key. This update disables module signing while building out-of-tree modules, thus, in the aforementioned case, kernel module loading works as expected.

BZ#[643815](#)¹⁴⁵

With this update, the upper limit of the *log_mtt_s_per_seg* variable was increased from five to seven, increasing the amount of memory that can be registered. As a result, the Mellanox driver (**mlx4**) can now use up to 64 GB of physical memory for RDMA (remote direct memory access). This provides better scalability for example when using the Mellanox adapter in NFS/RDMA, or on machines with a lot of physical memory.

BZ#[648408](#)¹⁴⁶

Due to a mix-up between *FMODE_* and *O_* flags, an NFSv4 client could get a *WRITE* lock on a file that another NFSv4 client already had a *READ* lock on. As a result, data could be corrupted. With this update, *FMODE_* and *O_* flags are properly handled and getting a *WRITE* lock fails in the aforementioned case.

BZ#[649436](#)¹⁴⁷

Booting Red Hat Enterprise Linux 6 debug kernel on a system with the Dell PowerEdge RAID Controller H700 adapter caused the *megaraid_sas* driver to reset the controller multiple times leading to a faulty controller state. On rebooting the system, the faulty controller state could

¹³⁶ <https://www.redhat.com/security/data/cve/CVE-2010-4077.html>

¹³⁷ <https://www.redhat.com/security/data/cve/CVE-2010-4079.html>

¹³⁸ <https://www.redhat.com/security/data/cve/CVE-2010-4080.html>

¹³⁹ <https://www.redhat.com/security/data/cve/CVE-2010-4081.html>

¹⁴⁰ <https://www.redhat.com/security/data/cve/CVE-2010-4082.html>

¹⁴¹ <https://www.redhat.com/security/data/cve/CVE-2010-4083.html>

¹⁴² <https://www.redhat.com/security/data/cve/CVE-2010-4158.html>

¹⁴³ <https://www.redhat.com/security/data/cve/CVE-2010-4525.html>

¹⁴⁴ https://bugzilla.redhat.com/show_bug.cgi?id=655122

¹⁴⁵ https://bugzilla.redhat.com/show_bug.cgi?id=643815

¹⁴⁶ https://bugzilla.redhat.com/show_bug.cgi?id=648408

¹⁴⁷ https://bugzilla.redhat.com/show_bug.cgi?id=649436

cause the firmware to detect an incorrect memory condition. This could be especially confusing since the message could be a faulty DIMM (Dual In-line Memory Module) condition prompting the administrator to replace the DIMMs. This occurred due to a leak in the `mfi_sg1` dma'ed frame when the firmware supported IEEE frames. The `mfi_sg1` would draw memory from the slab cache and any use of freed memory would result in incorrect pages being read in the ISR (Interrupt Service Routine). This caused the controller resets and the ensuing DIMM error condition. This update fixes the leak in `mfi_sg1` when the firmware supports IEEE frames. Faulty controller states and faulty DIMM conditions no longer occur.

BZ#[653900](https://bugzilla.redhat.com/show_bug.cgi?id=653900)¹⁴⁸

Running VDSM and performing an **lvextend** operation during an intensive Virtual Guest power up caused this operation to fail. Since **lvextend** was blocked, all components became non-responsive: **vgs** and **lvs** commands froze the session, Virtual Guests became *Paused* or *Not Responding*. This was caused due to a faulty use of a lock. With this update, performing an **lvextend** operation works as expected.

BZ#[651996](https://bugzilla.redhat.com/show_bug.cgi?id=651996)¹⁴⁹

Due to a faulty memory allocator, on Non-Uniform Memory Architecture (NUMA) platforms, an OOM (Out Of Memory) condition would occur when a user changed a cpuset's `/etc/dev/mems` file (list of memory nodes in that cpuset) even though the specified node had enough free memory. With this update, the memory allocator no longer causes an OOM condition when a node has enough free memory.

BZ#[653340](https://bugzilla.redhat.com/show_bug.cgi?id=653340)¹⁵⁰

When using a VIRT-IO (Virtual Input/Output) NIC (Network Interface Controller), its state was reported as *unknown* instead of its real state (*up* or *down*). This was due to the fact that the device could not report the state status. With this update, when a device is not capable of reporting the current state, it is assumed the state is *up* or the state is read from the config file.

BZ#[658879](https://bugzilla.redhat.com/show_bug.cgi?id=658879)¹⁵¹

A previously released patch fixed the external module compiling when using the full source tree, however, it was discovered it resulted in breaking the build in the kernel-devel only case. With this update, the patch has been fixed to avoid any external module compiling errors.

BZ#[647391](https://bugzilla.redhat.com/show_bug.cgi?id=647391)¹⁵²

Running certain workload tests on a NUMA (Non-Uniform Memory Architecture) system could cause kernel panic at `mm/migrate.c:113`. This was due to a false positive `BUG_ON`. With this update, the false positive `BUG_ON` has been removed.

BZ#[659611](https://bugzilla.redhat.com/show_bug.cgi?id=659611)¹⁵³

Updated partner qualification injecting target faults uncovered a flaw where the Emulex `lpfc` driver would incorrectly panic due to a null `pnode` dereference. This update addresses the issue and was tested successfully under the same test conditions without the panic occurring.

BZ#[660589](https://bugzilla.redhat.com/show_bug.cgi?id=660589)¹⁵⁴

Updated partner qualification injecting controller faults uncovered a flaw where the Emulex `lpfc` driver panicked during error handling. With this update, kernel panic no longer occurs.

¹⁴⁸ https://bugzilla.redhat.com/show_bug.cgi?id=653900

¹⁴⁹ https://bugzilla.redhat.com/show_bug.cgi?id=651996

¹⁵⁰ https://bugzilla.redhat.com/show_bug.cgi?id=653340

¹⁵¹ https://bugzilla.redhat.com/show_bug.cgi?id=658879

¹⁵² https://bugzilla.redhat.com/show_bug.cgi?id=647391

¹⁵³ https://bugzilla.redhat.com/show_bug.cgi?id=659611

¹⁵⁴ https://bugzilla.redhat.com/show_bug.cgi?id=660589

BZ#[660244](#)¹⁵⁵

Updated partner qualification injecting controller faults uncovered a flaw where Fibre Channel ports would go offline while testing with Emulex LPFC controllers due to a faulty LPFC heartbeat functionality. This update changes the default behavior of the LPFC heartbeat to *off*.

BZ#[660591](#)¹⁵⁶

When configuring an SIT (Simple Internet Transition) tunnel while a remote address is configured, kernel panic occurred, caused by an execution of a NULL *header_ops* pointer in the *neigh_update_hhs()* function. With this update, a check is introduced that makes sure the *header_ops* pointer is not of the value NULL, thus, kernel panic no longer occurs.

Users should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

B.1.6. RHSA-2010:0842: Important: kernel security and bug fix update



Important

This update has already been released (prior to the GA of this release) as the security errata [RHSA-2010:0842](#)¹⁵⁷

Updated kernel packages that fix multiple security issues and several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links after each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security fixes:

* Missing sanity checks in the Intel i915 driver in the Linux kernel could allow a local, unprivileged user to escalate their privileges. ([CVE-2010-2962](#)¹⁵⁸, Important)

* *compat_alloc_user_space()* in the Linux kernel 32/64-bit compatibility layer implementation was missing sanity checks. This function could be abused in other areas of the Linux kernel if its length argument can be controlled from user-space. On 64-bit systems, a local, unprivileged user could use this flaw to escalate their privileges. ([CVE-2010-3081](#)¹⁵⁹, Important)

* A buffer overflow flaw in *niu_get_ethtool_tcam_all()* in the *niu* Ethernet driver in the Linux kernel, could allow a local user to cause a denial of service or escalate their privileges. ([CVE-2010-3084](#)¹⁶⁰, Important)

¹⁵⁵ https://bugzilla.redhat.com/show_bug.cgi?id=660244

¹⁵⁶ https://bugzilla.redhat.com/show_bug.cgi?id=660591

¹⁵⁷ <https://rhn.redhat.com/errata/RHSA-2010-0842.html>

¹⁵⁸ <https://www.redhat.com/security/data/cve/CVE-2010-2962.html>

¹⁵⁹ <https://www.redhat.com/security/data/cve/CVE-2010-3081.html>

¹⁶⁰ <https://www.redhat.com/security/data/cve/CVE-2010-3084.html>

- * A flaw in the IA32 system call emulation provided in 64-bit Linux kernels could allow a local user to escalate their privileges. ([CVE-2010-3301](#)¹⁶¹, Important)
- * A flaw in `sctp_packet_config()` in the Linux kernel's Stream Control Transmission Protocol (SCTP) implementation could allow a remote attacker to cause a denial of service. ([CVE-2010-3432](#)¹⁶², Important)
- * A missing integer overflow check in `snd_ctl_new()` in the Linux kernel's sound subsystem could allow a local, unprivileged user on a 32-bit system to cause a denial of service or escalate their privileges. ([CVE-2010-3442](#)¹⁶³, Important)
- * A flaw was found in `sctp_auth_asoc_get_hmac()` in the Linux kernel's SCTP implementation. When iterating through the `hmac_ids` array, it did not reset the last id element if it was out of range. This could allow a remote attacker to cause a denial of service. ([CVE-2010-3705](#)¹⁶⁴, Important)
- * A function in the Linux kernel's Reliable Datagram Sockets (RDS) protocol implementation was missing sanity checks, which could allow a local, unprivileged user to escalate their privileges. ([CVE-2010-3904](#)¹⁶⁵, Important)
- * A flaw in `drm_ioctl()` in the Linux kernel's Direct Rendering Manager (DRM) implementation could allow a local, unprivileged user to cause an information leak. ([CVE-2010-2803](#)¹⁶⁶, Moderate)
- * It was found that wireless drivers might not always clear allocated buffers when handling a driver-specific IOCTL information request. A local user could trigger this flaw to cause an information leak. ([CVE-2010-2955](#)¹⁶⁷, Moderate)
- * A NULL pointer dereference flaw in `ftrace_regex_lseek()` in the Linux kernel's ftrace implementation could allow a local, unprivileged user to cause a denial of service. Note: The debugfs file system must be mounted locally to exploit this issue. It is not mounted by default. ([CVE-2010-3079](#)¹⁶⁸, Moderate)
- * A flaw in the Linux kernel's packet writing driver could be triggered via the **PKT_CTRL_CMD_STATUS** IOCTL request, possibly allowing a local, unprivileged user with access to `/dev/pktdvd/control` to cause an information leak. Note: By default, only users in the `cdrom` group have access to `/dev/pktdvd/control`. ([CVE-2010-3437](#)¹⁶⁹, Moderate)
- * A flaw was found in the way KVM (Kernel-based Virtual Machine) handled the reloading of `fs` and `gs` segment registers when they had invalid selectors. A privileged host user with access to `/dev/kvm` could use this flaw to crash the host. ([CVE-2010-3698](#)¹⁷⁰, Moderate)

Red Hat would like to thank Kees Cook for reporting CVE-2010-2962 and CVE-2010-2803; Ben Hawkes for reporting CVE-2010-3081 and CVE-2010-3301; Dan Rosenberg for reporting CVE-2010-3442, CVE-2010-3705, CVE-2010-3904, and CVE-2010-3437; and Robert Swiecki for reporting CVE-2010-3079.

Bug fixes:

¹⁶¹ <https://www.redhat.com/security/data/cve/CVE-2010-3301.html>

¹⁶² <https://www.redhat.com/security/data/cve/CVE-2010-3432.html>

¹⁶³ <https://www.redhat.com/security/data/cve/CVE-2010-3442.html>

¹⁶⁴ <https://www.redhat.com/security/data/cve/CVE-2010-3705.html>

¹⁶⁵ <https://www.redhat.com/security/data/cve/CVE-2010-3904.html>

¹⁶⁶ <https://www.redhat.com/security/data/cve/CVE-2010-2803.html>

¹⁶⁷ <https://www.redhat.com/security/data/cve/CVE-2010-2955.html>

¹⁶⁸ <https://www.redhat.com/security/data/cve/CVE-2010-3079.html>

¹⁶⁹ <https://www.redhat.com/security/data/cve/CVE-2010-3437.html>

¹⁷⁰ <https://www.redhat.com/security/data/cve/CVE-2010-3698.html>

BZ#632292¹⁷¹

When booting a Red Hat Enterprise Linux 5.5 kernel on a guest on an AMD host system running Red Hat Enterprise Linux 6, the guest kernel crashes due to an unsupported MSR (Model Specific Registers) read of the **MSR_K7_CLK_CTL** model. With this update, KVM support was added for the **MSR_K7_CLK_CTL** model specific register used in the AMD K7 CPU models, thus, the kernel crashes no longer occur.

BZ#633864¹⁷²

Previously, the s390 tape block driver crashed whenever it tried to switch the I/O scheduler. With this update, an official in-kernel API (`elevator_change()`) is used to switch the I/O scheduler safely, thus, the crashes no longer occurs.

BZ#633865¹⁷³

Previously, a kernel module not shipped by Red Hat was successfully loaded when the **FIPS** boot option was enabled. With this update, kernel self-integrity is improved by rejecting to load kernel modules which are not shipped by Red Hat when the **FIPS** boot option is enabled.

BZ#633964¹⁷⁴

A regression was discovered that caused kernel panic during the booting of any SGI UV100 and UV1000 system unless the **virtEFI** command line option was passed to the kernel by GRUB. With this update, the need for the **virtEFI** command line option is removed and the kernel will boots as expected without it.

BZ#633966¹⁷⁵

Previously, a Windows XP host experienced the stop error screen (i.e. the "Blue Screen Of Death" error) when booted with the CPU mode name. With this update, a Windows XP host no longer experiences the aforementioned error due to added KVM (Kernel-based Virtual Machine) support for the **MSR_EBC_FREQUENCY_ID** model specific register.

BZ#634973¹⁷⁶

Previously the cxgb3 (Chelsio Communications T3 10Gb Ethernet) adapter experienced parity errors. With this update, the parity errors are correctly detected and the cxgb3 adapter successfully recovers from them.

BZ#634984¹⁷⁷

Systems with an updated Video BIOS for the AMD RS880 would not properly boot with KMS (Kernel mode-setting) enabled. With this update, the Video BIOS boots successfully when KMS is enabled.

BZ#635951¹⁷⁸

The zfcpdump (kdump) kernel on IBM System z could not be debugged using the dump analysis tool **crash**, because the **vmLinux** file in the *kernel-kdump-debuginfo* RPM did not contain DWARF debug information. With this update, the **CONFIG_DEBUG_KERNEL** parameter is set to yes and the needed debug information is provided.

¹⁷¹ https://bugzilla.redhat.com/show_bug.cgi?id=632292

¹⁷² https://bugzilla.redhat.com/show_bug.cgi?id=633864

¹⁷³ https://bugzilla.redhat.com/show_bug.cgi?id=633865

¹⁷⁴ https://bugzilla.redhat.com/show_bug.cgi?id=633964

¹⁷⁵ https://bugzilla.redhat.com/show_bug.cgi?id=633966

¹⁷⁶ https://bugzilla.redhat.com/show_bug.cgi?id=634973

¹⁷⁷ https://bugzilla.redhat.com/show_bug.cgi?id=634984

¹⁷⁸ https://bugzilla.redhat.com/show_bug.cgi?id=635951

BZ#[636116](https://bugzilla.redhat.com/show_bug.cgi?id=636116)¹⁷⁹

Previously, **MADV_HUGEPAGE** was missing in the `include/asm-generic/mman-common.h` file which caused `madvise` to fail to utilize TPH. With this update, the `madvise` option was removed from `/sys/kernel/mm/redhat_transparent_hugepage/enabled` since **MADV_HUGEPAGE** was removed from the `madvise` system call.

BZ#[637087](https://bugzilla.redhat.com/show_bug.cgi?id=637087)¹⁸⁰

The kernel panicked when booting the `kdump` kernel on a `s390` system with an `initramfs` that contained an odd number of bytes. With this update, an `initramfs` with sufficient padding such that it contains an even number of bytes is generated, thus, the kernel no longer panics.

BZ#[638973](https://bugzilla.redhat.com/show_bug.cgi?id=638973)¹⁸¹

Previously, in order to install Snapshot 13, boot parameter `nomodeset xforcevesa` had to be added to the kernel command line, otherwise, the screen turned black and prevented the installation. With this update, the aforementioned boot parameter no longer has to be specified and the installation works as expected.

BZ#[639412](https://bugzilla.redhat.com/show_bug.cgi?id=639412)¹⁸²

Previously, a write request may have merged with a discard request. This could have posed a potential risk for 3rd party drivers which could possibly issue a discard without waiting properly. With this update, discarding of write block I/O requests by preventing merges of discard and write requests in one block I/O has been introduced, thus, resolving the possible risks.

BZ#[641258](https://bugzilla.redhat.com/show_bug.cgi?id=641258)¹⁸³, BZ#[644037](https://bugzilla.redhat.com/show_bug.cgi?id=644037)¹⁸⁴

The `fork()` system call led to an `rmap` walk finding the parent huge-pmd twice instead of once, thus causing a discrepancy between the `mapcount` and `page_mapcount` check, which could have led to erratic page counts for subpages. This fix ensures that the `rmap` walk is accurate when a process is forked, thus resolving the issue.

BZ#[641454](https://bugzilla.redhat.com/show_bug.cgi?id=641454)¹⁸⁵

Running a `fstress` test which issues various operations on a `ext4` filesystem when `usrquota` is enabled, the following JBD (Journaling Block Device) error was output in `/var/log/messages`:

```
JBD: Spotted dirty metadata buffer (dev = sda10, blocknr = 17635). There's a risk of
filesystem corruption in case of system crash.
```

With this update, by always journaling the quota file modification in an `ext4` file system the aforementioned message no longer appears in the logs.

BZ#[641455](https://bugzilla.redhat.com/show_bug.cgi?id=641455)¹⁸⁶

Previously, the destination MAC address validation was not checking for NPIV (N_Port ID Virtualization) addresses, which results in FCoE (Fibre Channel over Ethernet) frames being dropped. With this update, the destination MAC address check for FCoE frames has been modified so that multiple **N_port** IDs can be multiplexed on a single physical **N_port**.

¹⁷⁹ https://bugzilla.redhat.com/show_bug.cgi?id=636116

¹⁸⁰ https://bugzilla.redhat.com/show_bug.cgi?id=637087

¹⁸¹ https://bugzilla.redhat.com/show_bug.cgi?id=638973

¹⁸² https://bugzilla.redhat.com/show_bug.cgi?id=639412

¹⁸³ https://bugzilla.redhat.com/show_bug.cgi?id=641258

¹⁸⁴ https://bugzilla.redhat.com/show_bug.cgi?id=644037

¹⁸⁵ https://bugzilla.redhat.com/show_bug.cgi?id=641454

¹⁸⁶ https://bugzilla.redhat.com/show_bug.cgi?id=641455

BZ#641456¹⁸⁷

During an installation through Cisco NPV (N port virtualization) to Brocade, adding a LUN (Logical Unit Number) through **Add Advanced Target** did not work properly. This was caused by the faulty resending of FLOGI (Fabric Login) when a Fibre Channel switch in the NPV mode rejected requests with zero Destination ID. With this update, the LUN is seen and able to be selected for installation.

BZ#641457¹⁸⁸

Previously, timing issues could cause the FIP (FCoE Initialization Protocol) FLOGIs to timeout even if there were no problems. This caused the kernel to go into a non-FIP mode even though it should have been in the FIP mode. With this update, the timing issues no longer occur and the kernel no longer switches to the non-FIP mode when logging to the Fibre Channel Switch/Forwarder.

BZ#641458¹⁸⁹

Previously, the **vmstat** (virtual memory statistics) tool incorrectly reported the **disk I/O** as **swap-in** on ppc64 and other architectures that do not support the `TRANSPARENT_HUGEPAGE` configuration option in the kernel. With this update, the **vmstat** tool no longer reports incorrect statistics and works as expected.

BZ#641459¹⁹⁰

Previously, building under memory pressure with KSM (Kernel Shared Memory) caused KSM to collapse with an internal compiler error indicating an error in swapping. With this update, data corruption during swapping no longer occurs.

BZ#641460¹⁹¹

Occasionally, the `anon_vma` variable could contain the value `null` in the `page_address_in_vma` function and cause kernel panic. With this update, kernel panic no longer occurs.

BZ#641483¹⁹²

Previously, the `/proc/maps` file which is read by LVM2 (Logical Volume Manager 2) contained inconsistencies caused by LVM2 incorrectly deciding which memory to `mlock` and `munlock`. With this update, LVM2 correctly decides between the `mlock` and `munlock` operations and no longer causes inconsistencies.

BZ#641907¹⁹³

Systems that have an Emulex FC controller (with SLI-3 based firmware) installed could return a kernel panic during installation. With this update, kernel panic no longer occurs during installation.

BZ#642043¹⁹⁴

This update fixes the slow memory leak in the i915 module in DRM (Direct Rendering Manager) and GEM (Graphics Execution Manager).

¹⁸⁷ https://bugzilla.redhat.com/show_bug.cgi?id=641456

¹⁸⁸ https://bugzilla.redhat.com/show_bug.cgi?id=641457

¹⁸⁹ https://bugzilla.redhat.com/show_bug.cgi?id=641458

¹⁹⁰ https://bugzilla.redhat.com/show_bug.cgi?id=641459

¹⁹¹ https://bugzilla.redhat.com/show_bug.cgi?id=641460

¹⁹² https://bugzilla.redhat.com/show_bug.cgi?id=641483

¹⁹³ https://bugzilla.redhat.com/show_bug.cgi?id=641907

¹⁹⁴ https://bugzilla.redhat.com/show_bug.cgi?id=642043

BZ#642045¹⁹⁵

Previously, a race condition in the TTM (Translation Table Maps) module of the DRM (Direct Rendering Manager) between the object destruction thread and object eviction could result in a major loss of large objects reference counts. Consequently, this caused a major amount of memory leak. With this update, the race condition no longer occurs and any memory leaks are prevented.

BZ#642679¹⁹⁶

Previously, an operation such as `madvise(MADV_MERGEABLE)` may have split VMAs (Virtual Memory Area) without checking if any huge page had to be split into regular pages, leading to huge pages to be still mapped in VMA ranges that would not be large enough to fit huge pages. With this update, huge pages are checked whether they have been split when any VMA is being truncated.

BZ#642680¹⁹⁷

Previously, accounting of reclaimable inodes did not work correctly. When an inode was reclaimed it was only deleted from the per-AG (per Allocation Group) tree. Neither the counter was decreased, nor was the parent tree's AG entry untagged properly. This caused the system to hang indefinitely. With this update, the accounting of reclaimable inodes works properly and the system remains responsive.

BZ#644038¹⁹⁸

A race condition occurred when Xen was presented with an inconsistent page type resulting in the crash of the kernel. With this update, the race condition is prevented and kernel crashes no longer occur.

BZ#644636¹⁹⁹

Previously, Red Hat Enterprise Linux 6 enabled the `CONFIG_IMA` option in the kernel. This caused the kernel to track all inodes in the system in a radix tree, leading to a huge waste of memory. With this update, an optimized version of a tree (rbtree) is used and memory is no longer wasted.

BZ#644926²⁰⁰

Previously, calling the `elevator_change` function immediately after the `blk_init_queue` function resulted in a null pointer dereference. With this update, the null pointer dereference no longer occurs.

BZ#646994²⁰¹

When booting the latest Red Hat Enterprise Linux 6 kernel (-78.el6), the system hanged shortly after the booting. Access to the file system died and the console started outputting soft lockup messages from the TTM code. With this update, the aforementioned behavior no longer occurs and the system boots as expected.

Users should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

¹⁹⁵ https://bugzilla.redhat.com/show_bug.cgi?id=642045

¹⁹⁶ https://bugzilla.redhat.com/show_bug.cgi?id=642679

¹⁹⁷ https://bugzilla.redhat.com/show_bug.cgi?id=642680

¹⁹⁸ https://bugzilla.redhat.com/show_bug.cgi?id=644038

¹⁹⁹ https://bugzilla.redhat.com/show_bug.cgi?id=644636

²⁰⁰ https://bugzilla.redhat.com/show_bug.cgi?id=644926

²⁰¹ https://bugzilla.redhat.com/show_bug.cgi?id=646994

C. Revision History

Revision 1-5 **Thu May 19 2011**

Ryan Lerch r1erch@redhat.com

Removed Package Manifest data. Provided link to new Package Manifest document

Revision 1-5 **Tue Nov 16 2010**

Ryan Lerch r1erch@redhat.com

Fixed invalid links

Revision 1-0 **Wed Nov 10 2010**

Ryan Lerch r1erch@redhat.com

Initial Release of the Technical Notes

Index