



Red Hat Enterprise Linux 4

4.8 Release Notes

Release Notes for all Architectures

Edition 2

Last Updated: 2020-03-08

Red Hat Enterprise Linux 4 4.8 Release Notes

Release Notes for all Architectures

Edition 2

Ryan Lerch

Red Hat Engineering Content Services

rlersch@redhat.com

Legal Notice

Copyright © 2009 Red Hat.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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, the Red Hat logo, JBoss, OpenShift, 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.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This details the Release Notes for Red Hat Enterprise Linux 4.8

Table of Contents

1. RELEASE NOTES UPDATES	2
2. INTRODUCTION	3
2.1. Lifecycle	3
3. INSTALLATION-RELATED NOTES	4
3.1. All Architectures	4
3.2. ia64 Architectures	5
4. FEATURE UPDATES	6
4.1. All Architectures	6
5. KERNEL-RELATED UPDATES	7
5.1. All Architectures	7
5.2. x86-64 Architectures	11
5.3. s390x Architectures	11
6. DRIVER UPDATES	11
6.1. All Architectures	11
6.2. s390x Architectures	12
7. TECHNOLOGY PREVIEWS	13
8. RESOLVED ISSUES	13
8.1. All Architectures	13
9. KNOWN ISSUES	16
9.1. All Architectures	16
9.2. ia64 Architectures	21
A. REVISION HISTORY	22

1. RELEASE NOTES UPDATES

This section contains additional Release Notes and Updates to existing notes that were not included in the distribution version of the Red Hat Enterprise Linux 4.8 Release Notes.

Known Issue: [Bugzilla #488186](#)

Running Red Hat Enterprise Linux 4.8 on **Lenovo T61** notebooks may cause the system to hang during the boot process, displaying the following error message:

```
mtrr: v2.0 (20020519)
ACPI: Subsystem revision 20040816
ACPI: Found EC DT
```

To work around this issue, disable **ACPI (Advanced Configuration and Power Interface)** by adding the command **acpi=off** to the kernel boot parameters.

Known Issue: [Bugzilla #459785](#)

Japanese language JP106 keyboards will not function correctly when booting into Rescue Mode on **Red Hat Enterprise Linux 4.8**.

Known Issue: [Bugzilla #494022](#)

Updating all packages from **Red Hat Enterprise Linux 4.7** to **Red Hat Enterprise Linux 4.8** on multilib architectures may fail with dependency issues for the **openmpi-libs** package. To work around this issue, use the following commands to update the **compat-dapl** package before updating the remaining packages:

```
up2date compat-dapl
up2date -fu
```

Known Issue: [Bugzilla #443137](#)

In a typical *HA-RAID (High Availability RAID)* two-system configuration, two *SAS (Serial Attached SCSI)* adapters are plugged in to two systems that are connected to a shared SAS disk drawer. However, it is currently possible to set the **Preferred Dual Adapter State** attribute to **Primary** on both SAS adapters, which may trigger a race condition and cause infinite failover between the adapters. To prevent this error, if the **Preferred Dual Adapter State** attribute of one SAS adapters is set to **Primary**, ensure that the other is set to **None**.

Known Issue: [Bugzilla #499457](#)

As a result of **N_Port ID Virtualization (NPIV)** support added in **Red Hat Enterprise Linux 4.8** on s390x architectures, the **permanent_port_name sysfs** attribute is no longer included. This attribute was used (primarily for debugging purposes) to differentiate the use of NPIV **Logical Unit Numbers (LUNs)** from within Linux. In the absence of this attribute, system administrators should refer to the **Hardware Management Console / Support Element (HMC/SE)** to find the virtual port address on an NPIV-enabled system.

Known Issue: [Bugzilla #435300](#)

Known Issue: In previous versions of Red Hat Enterprise Linux 4, adding the line

-

```
selinux --permissive
```

to a kickstart file sets selinux to permissive mode. However, this line is currently ignored by the installer, leaving selinux set to the default mode: enforcing.

To set selinux to permissive mode during a kickstart installation, add the

```
setenforce 1
```

command to the

```
%pre
```

section of the kickstart file.

Alternatively, run

```
setenforce 1
```

after installation is complete.

Known Issue: [Bugzilla #455251](#)

In Red Hat Enterprise Linux 4, invoking the kernel system call **setpriority()** with a **which** parameter of type **PRIO_PROCESS** does not set the priority of child threads.

Recommendation: Firefox Restart

Red Hat strongly recommends restarting the Firefox browser after updating the **firefox** package. This will ensure that all Firefox updates take effect.

2. INTRODUCTION

The following topics are covered in this document:

- Installation-Related Notes
- Feature Updates
- Kernel-Related Updates
- Driver Updates
- Technology Previews
- Resolved Issues
- Known Issues

2.1. Lifecycle

The Red Hat Enterprise Linux 4 Life Cycle is available at:
<https://www.redhat.com/security/updates/errata/>

As previously announced, the release of Red Hat Enterprise Linux 4.8 will mark the beginning of Production 2 phase of the Red Hat Enterprise Linux 4. No new hardware enablement will be expected during this phase.

<https://www.redhat.com/archives/nahant-list/2008-July/msg00059.html>

Customers should note that their subscriptions provide access to all currently supported versions of Red Hat Enterprise Linux.

3. INSTALLATION-RELATED NOTES

The following section includes information specific to installation of Red Hat Enterprise Linux and the Anaconda installation program.



NOTE

When updating from one minor version of Red Hat Enterprise Linux 4 (such as 4.6 to 4.7) to Red Hat Enterprise Linux 4.8, it is recommended that you do so using Red Hat Network, whether through the hosted web user interface or Red Hat Network Satellite.

If you are upgrading a system with no available network connectivity, use the "Upgrade" functionality of **Anaconda**. However, note that **Anaconda** has limited abilities to handle issues such as dependencies on additional repositories or third-party applications. Further, **Anaconda** reports installation errors in a log file, not interactively.

As such, Red Hat recommends that when upgrading offline systems, you should test and verify the integrity of your upgrade configuration first. Be sure to carefully review the update log for errors before applying the upgrade to your production environment.

In-place upgrades between major versions of Red Hat Enterprise Linux (for example, upgrading from Red Hat Enterprise Linux 3 to Red Hat Enterprise Linux 4.8) is not supported. While the "Upgrade" option of **Anaconda** allows you to perform this, there is no guarantee that the upgrade will result in a working installation. In-place upgrades across major releases do not preserve all system settings, services, and custom configurations. For this reason, Red Hat strongly recommends that you perform a fresh installation when planning to upgrade between major versions.

3.1. All Architectures



IMPORTANT

If you are copying the contents of the Red Hat Enterprise Linux 4.8 CD-ROMs (in preparation for a network-based installation, for example) be sure you copy the CD-ROMs for the operating system *only*. Do not copy the Supplementary CD-ROM, or any of the layered product CD-ROMs, as this will overwrite files necessary for Anaconda's proper operation.

These CD-ROMs must be installed *after* Red Hat Enterprise Linux is installed.

Bugzilla #205295

The version of GRUB shipped with Red Hat Enterprise Linux 4 (and all updates) does not support software mirroring (RAID1). As such, if you install Red Hat Enterprise Linux 4 on a RAID1 partition, the bootloader will be installed in the first hard drive instead of the master boot record (MBR). This will render the system unbootable.

If you wish to install Red Hat Enterprise Linux 4 on a RAID1 partition, you should clear any pre-existing bootloader from the MBR first.

Bugzilla #222958

When installing Red Hat Enterprise Linux 4 in Text Mode on systems that use flat-panel monitors and some *ATI* cards, the screen area may appear shifted. When this occurs, some areas of the screen will be obscured.

If this occurs, perform the installation with the parameter **linux nofb**.

Bugzilla #445835

When upgrading from Red Hat Enterprise Linux 4.6 to this release, **minilogd** may log several SELinux denials. These error logs are harmless, and can be safely ignored.

Bugzilla #430476

Previously, in the **Anaconda** kickstart documentation (located at: `/usr/share/doc/anaconda-<anaconda-version>/kickstart-docs.txt`), the section detailing the **--driveorder** option in a kickstart file stated:

Specify which drive is first in the BIOS boot order.

However, the **--driveorder** option actually requires a list of all drives on the system, with the first boot device appearing first in the list. With this update, the documentation has been clarified and now reads:

Specify which drive is first in the BIOS boot order.
The ordered list must include all the drives in the system.

When using the **--driveorder** option in a kickstart file The ordered list must include all the drives in the system.

3.2. ia64 Architectures

Bugzilla #163910

In this update, the 64-bit Intel Itanium2 architecture includes runtime support for 32-bit applications through the use of Intel's IA-32 Execution Layer.

The IA-32 Execution Layer is provided on the Extras disc for the Intel Itanium2 architecture. In addition, a set of 32-bit libraries and applications are provided on a separate 32-bit Compatibility Layer disc. The IA-32 Execution Layer and 32-bit compatibility packages together provide a runtime environment for 32-bit applications on the 64-bit native distribution.

To install the IA-32 Execution Layer and required 32-bit compatibility packages, follow these steps:

1. Install Red Hat Enterprise Linux 4.8 for the Intel Itanium2 Architecture.
2. Insert the Red Hat Enterprise Linux 4 Extras CD, which contains the **ia32el** package.
3. After the system has mounted the CD, change to the directory containing the Extras packages. For example:

```
cd /media/cdrom/RedHat/RPMS/
```

4. Install the **ia32el** and **ksh** packages:

```
rpm -Uvh ia32el-<version>.ia64.rpm ksh-<version>.ia64.rpm
```

where <version> is the respective versions of the **ia32el** and **ksh** packages to be installed.

5. Eject the Extras CD:

```
eject /media/cdrom
```

6. To verify the installation of the 32-bit compatibility layer and libraries after installation, confirm if the **/emul** directory exists and contains files.

7. To verify that the 32-bit compatibility mode is in effect, type the following in a shell prompt:

```
service ia32el status
```

8. At this point you can install compatibility libraries by inserting the 32-bit Compatibility Layer disc. You may choose to install all of the packages available on the disc or choose the particular packages required in order to provide runtime support for your 32-bit applications.

4. FEATURE UPDATES

4.1. All Architectures

Bugzilla #469924

Systemtap is now a fully supported feature in Red Hat Enterprise Linux 4. **systemtap** provides a free software (GPL) infrastructure to simplify the gathering of information about the running Linux system. This assists diagnosis of a performance or functional problem. With the help of **systemtap**, developers no longer need to go through the tedious and disruptive sequence of instrument, recompile, install, and reboot that may be otherwise required to collect data.

Note that some features of **systemtap** for newer Red Hat Enterprise Linux or Linux systems will not work on Red Hat Enterprise Linux 4 due to missing kernel features. The absence of the kernel **utrace** precludes support for any user-space probing.

Bugzilla #459041

dmidecode gives information about BIOSes and motherboard revisions. The version of kernel-utils supplied with this advisory updates **dmidecode** from version 2.2 to version 2.9. This version identifies newer processors, PCI-express slots and devices, and blade chassis. It also offers enhanced support for the SMBIOS v2.6 specification.

Bugzilla #453642

a new version of **kernel-utils** is included in this release, updating the Intel microcode file to version 20080910, to support newer Intel processors.

Bugzilla #447979

smartmontools has been extended to support newer CCISS controllers found in newer HP ProLiant hardware.

Bugzilla #460904

The **Samba** package has been rebased to the upstream version 3.0.33. The 3.0.x version series is a bugfix only branch of the **Samba** code base. By rebasing to 3.0.33 we will include a number of important bug fixes and security fixes. No new features will be added by this rebase.

For more information on the upstream fixes provided by this rebase, refer to the Samba Release Notes: <http://samba.org/samba/history/samba-3.0.33.html>

Bugzilla #454833

ipmitool has been updated to the upstream version 1.8.11, which provides several bug fixes and enhancements over the previous release, including the following:

- Documentation update
- Bugfixes for SDR/FRU, SOL and many others
- New commands and options

Please note that behaviour of the **-K** command line switch has changed from **prompt for Kg key** to **read Kg key from environment variable**. The **-Y** flag now behaves as the **-K** did prior to this update.

5. KERNEL-RELATED UPDATES

5.1. All Architectures

Bugzilla #467714

The **ibmphp** module is not safe to unload. Previously, the mechanism that prevented the **ibmphp** module from unloading was insufficient, and eventually triggered a bug halt. With this update, the method to prevent this module from unloading has been improved, preventing the bug halt. However, attempting to unload the module may produce a warning in the message log, indicating that the module is not safe to unload. This warning can be safely ignored.

Bugzilla #461564

With this update, physical memory will be limited to 64GB for 32-bit x86 kernels running on systems with more than 64GB. The kernel splits memory into 2 separate regions: Lowmem and Highmem. Lowmem is mapped into the kernel address space at all times. Highmem, however, is mapped into a kernel virtual window a page at a time as needed. If memory I/Os are allowed to exceed 64GB, the `mem_map` (also known as the page array) size can approach or even exceed the size of Lowmem. If this happens, the kernel panics during boot or starts prematurely. In the latter case, the kernel fails to allocate kernel memory after booting and either panics or hangs.

Bugzilla #246233

Previously, if a user pressed the arrow keys continuously on a Hardware Virtual Machine (HVM) an interrupt race condition between the hardware interrupt and timer interrupt was encountered. As a result, the keyboard driver reported unknown keycode events. With this update, the i8042 polling timer has been removed, which resolves this issue.

Bugzilla #435705

With this update, the `diskdump` utility (which provides the ability to create and collect vmcore Kernel dumps) is now supported for use with the **sata_svw** driver.

Bugzilla #439043

With this update, the "swap_token_timeout" parameter has been added to `/proc/sys/vm`.

This file contains valid hold time of swap out protection token. The Linux Virtual Memory (VM) subsystem has a token based thrashing control mechanism and uses the token to prevent unnecessary page faults in thrashing situation. The unit of the value is in `second`. The value would be useful to tune thrashing behavior. Setting it to 0 will disable the swap token mechanism.

Bugzilla #439431

Previously, when a NFSv4 (Network File System Version 4) client encountered issues while processing a directory using `readdir()`, an error for the entire `readdir()` call was returned. With this update, the `fatrr4_rdatrr_error` flag is now set when `readdir()` is called, instructing the server to continue on and only report an error on the specific directory entry that was causing the issue.

Bugzilla #443655

Previously, the NFS (Network File System) client was not handling malformed replies from the `readdir()` function. Consequently, the reply from the server would indicate that the call to the `readdir()` function was successful, but the reply would contain no entries. With this update, the `readdir()` reply parsing logic has been changed, such that when a malformed reply is received, the client returns an EIO error.

Bugzilla #448076

The RPC client stores the result of a portmap call at a place in memory that can be freed and reallocated under the right circumstances. However, under some circumstances, the result of the portmap call was freed from memory too early, which may have resulted in memory corruption. With this update, reference counting has been added to the memory location where the portmap result is stored, and will only free it after it has been used.

Bugzilla #450743

Under some circumstances, the allocation of some data structures for RPC calls may have been blocked when the system memory was low. Consequently, deadlock may have been encountered under heavy memory pressure when there were a large number of NFS pages awaiting writeback. With this update, the allocation of these data structures is now non-blocking, which resolves this issue.

Bugzilla #451088

Previously, degraded performance may have been encountered when writing to a LVM mirrored volume synchronously (using the `O_SYNC` flag). Consequently, every write I/O to a mirrored volume was delayed by 3ms, resulting in the mirrored volume being approximately 5-10 times slower than a linear volume. With this update, I/O queue unplugging has been added to the `dm-raid1` driver, and the performance of mirrored volumes has been improved to be comparable with that of linear volumes.

Bugzilla #476997

A new tuning parameter has been added to allow system administrators to change the max number of modified pages `kupdate` writes to disk per iteration each time it runs. This new tunable (`/proc/sys/vm/max_writeback_pages`) defaults to a value of 1024 (4MB) so that a maximum of 1024 pages get written out by each iteration of `kupdate`. Increasing this value alters how aggressively `kupdate` flushes modified pages and decreases the potential amount of data loss if the system crashes between `kupdate` runs. However, increasing the `max_writeback_pages` value may have negative performance consequences on systems that are sensitive to I/O loads.

Bugzilla #456911

A new allowable value has been added to the `/proc/sys/kernel/wake_balance` tunable parameter. Setting `wake_balance` to a value of 2 will instruct the scheduler to run the thread on any available CPU rather than scheduling it on the optimal CPU. Setting this kernel parameter to 2 will force the scheduler to reduce the overall latency even at the cost of total system throughput.

Bugzilla #475715

When checking a directory tree, the kernel module could, in some circumstances, incorrectly decide the tree was not busy. An active offset mount with an open file handle being used for expires caused the file handle to not count toward the busyness check. This resulted in mount requests being made for already mounted offsets. With this update, the kernel module check has been corrected and incorrect mount requests are no longer generated.

Bugzilla #453470

During system initialization, the CPU vendor was detected after the initialization of the Advanced Programmable Interrupt Controllers (APICs). Consequently, on x86_64 AMD systems with more than 8 cores, APIC clustered mode was used, resulting in suboptimal system performance. With this update, the CPU vendor is now queried prior to initializing the APICs, resulting in APIC physical flat mode being used by default, which resolves this issue.

Bugzilla #462459

The **Common Internet File System (CIFS)** code has been updated in Red Hat Enterprise Linux 4.8, fixing a number of bugs that had been repaired in upstream, including the following change:

Previously, when mounting a server without Unix extensions, it was possible to change the mode of a file. However, this mode change could not be permanently stored, and may have changed back to the original mode at any time. With this update, the mode of the file cannot be temporarily changed by default; `chmod()` calls will return success, but have no effect. A new mount option, `dynperm` needs to be used if the old behavior is required.

Bugzilla #451819

Previously, in the kernel, there was a race condition may have been encountered between `dio_bio_end_aio()` and `dio_await_one()`. This may have lead to a situation where direct I/O is left waiting indefinitely on an I/O process that has already completed. With this update, these reference counting operations are now locked so that the submission and completion paths see a unified state, which resolves this issue.

Bugzilla #249775

Previously, upgrading a fully virtualized guest system from Red Hat Enterprise Linux 4.6 (with the `kmod-xenpv` package installed) to newer versions of Red Hat Enterprise Linux 4 resulted in an improper module dependency between the built-in kernel modules: `xen-vbd.ko` & `xen-vnif.ko` and the older `xen-platform-pci.ko` module. Consequently, file systems mounted via the `xen-vbd.ko` block driver, and guest networking using the `xen-vnif.ko` network driver would fail.

In Red Hat Enterprise Linux 4.7, the functionality in the `xen-platform-pci.ko` module was built-in to the kernel. However, when a formally loadable kernel module becomes a part of the kernel, the symbol dependency check for existing loadable modules is not accounted for in the module-init-tools correctly. With this update, the `xen-platform-pci.ko` functionality has been removed from the built-in kernel and placed back into a loadable module, allowing the module-init-tools to check and create the proper dependencies during a kernel upgrade.

Bugzilla #463897

Previously, attempting to mount disks or partitions in a 32-bit Red Hat Enterprise Linux 4.6 fully

virtualized guest using the paravirtualized block driver (**xen-vbd.ko**) on a 64-bit host would fail. With this update, the block front driver (**block.c**) has been updated to inform the block back driver that the guest is using the 32-bit protocol, which resolves this issue.

Bugzilla #460984

Previously, installing the **pv-on-hvm** drivers on a bare-metal kernel automatically created the **/proc/xen** directory. Consequently, applications that verify if the system is running a virtualized kernel by checking for the existence of the **/proc/xen** directory may have incorrectly assumed that the virtualized kernel is being used. With this update, the **pv-on-hvm** drivers no longer create the **/proc/xen** directory, which resolves this issue.

Bugzilla #455756

Previously, paravirtualized guests could only have a maximum of 16 disk devices. In this update, this limit has been increased to a maximum of 256 disk devices.

Bugzilla #523930

In some circumstances, write operations to a particular TTY device opened by more than one user (eg, one opened it as **/dev/console** and the other opened it as **/dev/ttyS0**) were blocked. If one user opened the TTY terminal without setting the **O_NONBLOCK** flag, this user's write operations were suspended if the output buffer was full or if a **STOP** (Ctrl-S) signal was sent. As well, because the **O_NONBLOCK** flag was not respected, write operations for user terminals opened with the **O_NONBLOCK** flag set were also blocked. This update re-implements TTY locks, ensuring **O_NONBLOCK** works as expected, even if a **STOP** signal is sent from another terminal.

Bugzilla #519692

Previously, the **get_random_int()** function returned the same number until the jiffies counter (which ticks at a clock interrupt frequency) or process ID (PID) changed, making it possible to predict the random numbers. This may have weakened the ASLR security feature. With this update, **get_random_int()** is more random and no longer uses a common seed value. This reduces the possibility of predicting the values **get_random_int()** returns.

Bugzilla #518707

ib_mthca, the driver for Host Channel Adapter (HCA) cards based on the Mellanox Technologies MT25408 InfiniHost III Lx HCA integrated circuit device, uses **kmalloc()** to allocate large bitmaps. This ensures allocated memory is a contiguous physical block, as is required by DMA devices such as these HCA cards.

Previously, the largest allowed **kmalloc()** was a 128kB page. If **ib_mthca** was set to allocate more than 128kB (for example, by setting the **num_mutt** option to "num_mutt=2097152", causing **kmalloc()** to allocate 256kB) the driver failed to load, returning the message

```
Failed to initialize memory region table, aborting.
```

This update alters the allocation methods of the **ib_mthca** driver. When **mthca_buddy_init()** wants more than a page, memory is allocated directly from the page allocator, rather than using **kmalloc()**. It is now possible to pin large amounts of memory for use by the **ib_mthca** driver by increasing the values assigned to **num_mutt** and **num_mtt**.

Bugzilla #519446

Previously, there were some instances in the kernel where the **__ptrace_unlink()** function (part of the **ptrace** system call) used **REMOVE_LINKS** and **SET_LINKS**, rather than **add_parent** and **remove_parent**, while changing the parent of a process. This approach could abuse the global

process list and, as a consequence, create deadlocked and unkillable processes in some circumstances. With this update, `__ptrace_unlink()` now uses `add_parent` and `remove_parent` in every instance, ensuring that deadlocked and unkillable processes cannot be created.



NOTE

Unkillable or deadlocked processes created by this bug had no effect on system availability.

5.2. x86-64 Architectures

Bugzilla #437881

Previously, there was a missing sign extension in the `x86_64 ptrace` code that may have caused `gdb` to fail on the `x86_64` architecture when debugging an `i386` application. With this update, the missing sign extension is now correctly extended, which resolves this issue.

5.3. s390x Architectures

Bugzilla #249775

On **Red Hat Enterprise Linux 4.8**, N_Port ID Virtualization (NPIV) for System z guests using zFCP is now enabled. NPIV allows a Fibre Channel HBA to log in multiple times to a Fibre Channel fabric using a single physical port (N_Port). With this functionality, a Storage Area Network (SAN) administrator can assign one or more logical unit numbers (LUNs) to a particular System z guest, making that LUN inaccessible to others. For further information, see "Introducing N_Port Identifier Virtualization for IBM System z9, REDP-4125" available at <http://www.redbooks.ibm.com/abstracts/redp4125.html>

6. DRIVER UPDATES

6.1. All Architectures

Bugzilla #452846

The Intel® High Definition Audio (HDA) driver in ALSA has been updated. This update improves audio support for newer hardware with HDA integrated audio.

Bugzilla #479408

Previously, network devices using the `forcedeth` driver may have stopped responding while doing `rpc` command from multiple clients. With this update, the `forcedeth` driver has been updated, which resolves this issue.

Bugzilla #441707

Previously, the Automatic Direct Memory Access (ADMA) mode was enabled by default in the `sata_nv` driver. Consequently, device errors and timeouts may have been encountered with some devices that utilize the `sata_nv` driver. With this update, ADMA mode is now disabled by default, which resolves this issue.

Bugzilla #446215

The drivers for `virtio`, the platform for I/O virtualization in KVM, has been backported to Red Hat

Enterprise Linux 4.8 from Linux Kernel 2.6.27. These drivers will enable KVM guests to achieve higher levels of I/O performance. Various user space components such as: **anaconda**, **kudzu**, **lvm**, **selinux** and **mkinitrd** have also been updated to support **virtio** devices.

Bugzilla #451966

The r8169 driver has been updated to provide support for newer network chipsets. With this update, all variants of RTL810x/RTL8168(9) are now supported in Red Hat Enterprise Linux 4.8.

Bugzilla #452163

The mptsas driver has been updated to version 3.12.29.00. This update includes bug fixes and enables the following new features:

- Dual Port support.
- SAS chip Power Management.

Bugzilla #452271

The lpfc driver has been updated to version number to 8.0.16.46. This update applies several bug fixes and enhancements, including:

- support for FCoE LP21000 HBAs
- support for HBAnyware 4.0

Bugzilla #455297

The megaraid_sas driver for SAS based RAID controllers has been updated to version 4.01-RH1. Several bug fixes and improvements are applied by this update, including:

- Added support for the LSI Generation 2 Controllers (0078, 0079)
- Added a command to shutdown DCMD in the shutdown routine to improve firmware shutdown.
- A bug that caused unexpected interrupts in the hardware Linux driver has been fixed.

Bugzilla #454838

The eHEA ethernet device driver for **IBM eServer System P** has been updated to version 0078-08.

Bugzilla #490503

The **EHCA infinband** device driver will not be supported for Red Hat Enterprise Linux 4.8 and all future Red Hat Enterprise Linux 4 releases.

6.2. s390x Architectures

Bugzilla #448777

Systems using zFCP for access to SCSI disks on Red Hat Enterprise Linux 4 require a hardware fibre channel switch to be connected between the mainframe and disk storage. This update enables point-to-point connections, which are fibre connections directly from the mainframe to the disk storage. While connection to a fibre channel switch is still supported, it is no longer required.

7. TECHNOLOGY PREVIEWS

Technology Preview features are currently *not* supported under Red Hat Enterprise Linux 4.8 subscription services, may not be functionally complete, and are generally not suitable for production use. However, these features are included as a customer convenience and to provide the feature with wider exposure.

Customers may find these features useful in a non-production environment. Customers are also free to provide feedback and functionality suggestions for a technology preview feature before it becomes fully supported. Erratas will be provided for high-severity security issues.

During the development of a technology preview feature, additional components may become available to the public for testing. It is the intention of Red Hat to fully support technology preview features in a future release.

For more information on the scope of Technology Previews in Red Hat Enterprise Linux, please view the [Technology Preview Features Support Scope](#) page on the Red Hat website.

OpenOffice 2.0

OpenOffice 2.0 is now included in this release as a Technology Preview. This suite features several improvements, including ODF and PDF functionalities, support for digital signatures and greater compatibility with open suites in terms of format and interface. In addition to this, the *OpenOffice 2.0* spreadsheet has enhanced pivot table support, and can now handle up to 65,000 rows.

For more information about *OpenOffice 2.0*, please refer to http://www.openoffice.org/dev_docs/features/2.0/index.html.

8. RESOLVED ISSUES

8.1. All Architectures

Bugzilla #452919

Previously, if the **Red Hat Network applet** was used to re-register the client to a different **Red Hat Satellite Server**, the applet would continue to show updates that had been available on the previous server, even if they were not available on the current server. The **/etc/sysconfig/rhn/rhn-applet** would not change to reflect the details of the new server. The version of the applet provided with this update associates a cache of updates with a server url, and therefore ensure that the updates displayed to the user are actually available. This version can also detect when its configuration file has changed. If such a change is detected, the applet will automatically reload the configuration variables and create new server connections.

Bugzilla #454690

On some *SGI Altix* systems that feature the IOC4 multi-function device, you may encounter problems when using attached IDE devices (such as CD-ROM drives). This is caused by a bug in the **sgiioc4** IDE driver, which prevents some devices from being detected properly on system boot.

You can work around this bug by manually loading the driver, which in turn allows attached IDE devices to be detected properly. To do so, run the following command as root:

```
/sbin/modprobe sgiioc4
```

Bugzilla #454690

sysreport.legacy used **\$HOME** as its root directory. In case this environment variable did not exist or the directory it referred to was not writable, **sysreport.legacy** could not generate its report and would exit with the message **Cannot make temp dir**. **Sysreport.legacy** now uses a randomly created directory as its root directory and therefore can generate a report even on a system without a usable **\$HOME**.

Bugzilla #476767

The **automount** daemon used fixed size buffer of 128 bytes long to receive information from the **SIOCGIFCONF ioctl** about local interfaces when testing for the proximity of a host corresponding to a given mount. Since the details of each interface are 40 bytes long, the daemon could receive information on no more than three local interfaces. If the host corresponding to the mount had an address that was local but did not correspond to one of the three interfaces the proximity would be classified incorrectly.

The **automount** daemon now dynamically allocates a buffer, ensuring that it is large enough to contain information on all interfaces on the system providing the ability to correctly detect proximity of a host given for an NFS mount.

Bugzilla #465237

Automount map entries that refer to multiple hosts in the mount location (replicated mount), the **automount** daemon probes a list of remote hosts for their proximity and NFS version. If hosts fail to respond, they are removed from the list. If no remote hosts reply at all, the list may become empty. Previously, the daemon did not check if the list was empty following the initial probe which would lead to a segmentation fault (by dereferencing a NULL pointer). This check has been added.

Bugzilla #444942

the **ttfonts-zh_CN** package formerly included the Zhong Yi Song TrueType font. The copyright in this font belongs to Beijing Zhong Yi Electronics Co., which has licensed Red Hat Inc. to distribute the font only in products and software under the Red Hat name. The inclusion of this font in **ttfonts-zh_CN** would therefore preclude Red Hat from freely distributing this package. The Zhong Yi Song TrueType font is still available to Red Hat customers via the Red Hat Network and the Supplementary CD in the **fonts-chinese-zysong** package.

Bugzilla #457228

multipathd crashed with a status of with a **multipathd dead but pid file exists** when multipath was configured for 1024 or more paths, because it was unable to open a file descriptor for each path. This may also have caused **error calling out /sbin/mpath_prio_ontap /dev/[device]** errors. Now, a new **multipath.conf** parameter, **max_fds**, allows end-users to set the maximum number of file descriptors that the **multipathd** process can have open, or to use **max** to set the number to the system maximum. Setting **max_fds** to either a sufficiently high number or to **max** avoids this crash in **multipathd**.

Bugzilla #457552

Previously, when using the **accraid** driver with an **Adaptec 2120S** or **Adaptec 2200S** controller, the system may have failed to bootup, returning the error: **aac_srb:aac_fib_send failed with status 8195**. With this update, the **accraid** driver has been updated, which resolves this issue.

Bugzilla #453150

SOS is a set of tools that gathers information about a system's hardware and current configuration. The information can then be used for diagnostic purposes and debugging.

With this update, the reports generated by `sosreport` now include five types of information that were not previously collected:

- the content of `/var/log/cron*` and the output of `crontab -l` to show what was running at the time that the problem occurred.
- partition information from `parted` instead of what was previously collected from `fdisk`, since `parted` can collect partition information in situations where `fdisk` cannot (for example, GUID partitions).
- output from `dumpe2fs -l`.
- the content of `/etc/inittab`.
- output from `"/sbin/service --status-all"` to show the current status of services. Previously, only their settings at boot time were collected (from `"chkconfig --list"`).

Bugzilla #453999

automount uses **umount(8)** when expiring mounts and **umount(8)** can wait indefinitely for a server to respond. This can lead to the expire being blocked causing mounts not to be expired for a long period of time in the same **/usr/sbin/automount** process (that is, the mount that the given automount process is managing). Consequently, if a server was unreachable, then automount would not unmount any expired mounts, even on the servers that are responding. Systems can then be left with a large number of mounts that can be expired but are not. Automount now includes a command line option to specify a time for automount to wait before giving up and moving on to remaining mounts. Expired mounts can therefore be unmounted even if some servers do not respond.

Bugzilla #479016

The `netpbm` package has been updated to fix the following bugs:

- Several utilities shipped with `netpbm` did not accept files from standard input even though this method was in accordance with documentation. With this update, this issue has been resolved.
- Several utilities shipped with `netpbm` may have crashed during processing of image files. With this update, this issue has been resolved.

Bugzilla #490104

the **ICQ** Internet message protocol servers recently changed and now require clients to use a newer version of the ICQ protocol. Logging in to ICQ with **Pidgin 2.5.2** (the version previously shipped with Red Hat Enterprise Linux 4) failed with an error message as a result. With this update, **Pidgin** has been updated to version 2.5.5, which resolves this issue.

Bugzilla #479692

Previously, the Red Hat Knowledgebase article documenting Fibre Channel rescan in Red Hat Enterprise Linux 4 was not accurate. This procedure has now been updated, and can be viewed at: <http://kbase.redhat.com/faq/docs/DOC-3942>

Bugzilla #422371

After successfully connecting to an SSH server, the server may return a text based banner back to the SSH client. Consequently, if **gftp** (a graphical ftp client) attempted to connect (via SFTP) to an SSH server that returns a banner, **gftp** would interpret the banner as an error, and close the connection. With this update, **gftp** has been updated to version 2.0.18, allowing connections to servers with banners.

Bugzilla #452257

When uploading a single file to a NFS directory, the timestamp indicating the modification and access times of the file may not have been recorded correctly. With this update, the timestamp is now always updated, which resolves this issue.

Bugzilla #453033

The probing code in **kudzu** for PCI devices would not properly find some modules that work by binding to specific PCI classes, notably, the `sgioc4` driver on SGI Altix systems. Without these modules loaded, the system would not detect devices that depended on the driver. A new version of the probing code is included in this updated package, which is able to successfully find the affected modules.

9. KNOWN ISSUES

9.1. All Architectures

Bugzilla #484117

The Logical Volume Manager in Red Hat Enterprise Linux 4.8 reports file descriptor leaks, resulting in the following error returned to the installation output:

```
File descriptor NUM (socket:XXXX) leaked on lvm invocation.
```

This message can be safely ignored.

Bugzilla #468097

When installing Red Hat Enterprise Linux 4 through an Network File System (NFS) server, the installer is unable to correctly close the NFS mount points. This might cause the NFS server to misbehave. In these cases Red Hat suggests the use of an HTTP server for installations.

Bugzilla #468097

On systems where the BIOS is able to do both legacy (**acpiphp**) and native (**pciehp**) PCI hotplugging, it is necessary for the administrator to choose a preferred method and explicitly prevent Red Hat Enterprise Linux 4 from loading the module for the undesired method. This is done by blacklisting the undesired module in **/etc/modprobe.conf**.

Bugzilla #451164

Hardware testing for the Mellanox MT25204 has revealed that an internal error occurs under certain high-load conditions. When the `ib_mthca` driver reports a catastrophic error on this hardware, it is usually related to an insufficient completion queue depth relative to the number of outstanding work requests generated by the user application.

Although the driver will reset the hardware and recover from such an event, all existing connections at the time of the error will be lost. This generally results in a segmentation fault in the user application. Further, if `opensm` is running at the time the error occurs, then you need to manually restart it in order to resume proper operation.

Bugzilla #443795

A bug in previous versions of **openmpi** and **lam** may prevent you from upgrading these packages. This same bug may cause **up2date** to fail when upgrading all packages.

This bug manifests in the following error when attempting to upgrade **openmpi** or **lam**:

```
error: %preun(openmpi-[version]) scriptlet failed, exit status 2
```

This bug also manifests in the following error (logged in `/var/log/up2date`) when attempting to upgrade all packages through **up2date**:

```
up2date Failed running rpm transaction - %pre %pro failure ?.
```

As such, you need to manually remove older versions of **openmpi** and **lam** first in order to avoid these errors. To do so, use the following **rpm** command:

```
rpm -qa | grep '^openmpi-|^lam-' | xargs rpm -e --noscripts --allmatches
```

Bugzilla #430494

When a LUN is deleted on a configured storage system, the change is not reflected on the host. In such cases, **lvm** commands will hang indefinitely when **dm-multipath** is used, as the LUN has now become *stale*.

To work around this, delete all device and **mpath** link entries in `/etc/lvm/.cache` specific to the stale LUN. To find out what these entries are, run the following command:

```
ls -l /dev/mpath | grep <stale LUN>
```

For example, if `<stale LUN>` is `3600d0230003414f30000203a7bc41a00`, the following results may appear:

```
lrwxrwxrwx 1 root root 7 Aug  2 10:33 /3600d0230003414f30000203a7bc41a00 -> ../dm-4
lrwxrwx--rwx 1 root root 7 Aug  2 10:33 /3600d0230003414f30000203a7bc41a00p1 -> ../dm-5
```

This means that `3600d0230003414f30000203a7bc41a00` is mapped to two **mpath** links: **dm-4** and **dm-5**.

As such, the following lines should be deleted from `/etc/lvm/.cache`:

```
dev/dm-4
dev/dm-5
dev/mapper/3600d0230003414f30000203a7bc41a00
dev/mapper/3600d0230003414f30000203a7bc41a00p1
dev/mpath/3600d0230003414f30000203a7bc41a00
dev/mpath/3600d0230003414f30000203a7bc41a00p1
```

Bugzilla #195685

If you need to use the **hp_sw** kernel module, install the updated **device-mapper-multipath** package.

You also need to properly configure the *HP* array to correctly use active/passive mode and recognize connections from a Linux machine. To do this, perform the following steps:

1. Determine what the world wide port name (WWPN) of each connection is by using **show connections**. Below is a sample output of **show connections** on an *HP MSA1000* array with two connections:

```
Connection Name: <Unknown>
```

```

Host WWNN = 200100E0-8B3C0A65
Host WWPN = 210100E0-8B3C0A65
Profile Name = Default
Unit Offset = 0
Controller 2 Port 1 Status = Online

```

```

Connection Name: <Unknown>
Host WWNN = 200000E0-8B1C0A65
Host WWPN = 210000E0-8B1C0A65
Profile Name = Default
Unit Offset = 0
Controller 1 Port 1 Status = Online

```

2. Configure each connection properly using the following command:

```
add connection [connection name] WWPN=[WWPN ID] profile=Linux OFFSET=[unit offset]
```

Note that **[connection name]** can be set arbitrarily.

Using the given example, the proper commands should be:

```
add connection foo-p2 WWPN=210000E0-8B1C0A65 profile=Linux OFFSET=0
```

```
add connection foo-p1 WWPN=210100E0-8B3C0A65 profile=Linux OFFSET=0
```

3. Run **show connections** again to verify that each connection is properly configured. As per the given example, the correct configuration should be:

```

Connection Name: foo-p2
Host WWNN = 200000E0-8B1C0A65
Host WWPN = 210000E0-8B1C0A65
Profile Name = Linux
Unit Offset = 0
Controller 1 Port 1 Status = Online

```

```

Connection Name: foo-p1
Host WWNN = 200100E0-8B3C0A65
Host WWPN = 210100E0-8B3C0A65
Profile Name = Linux
Unit Offset = 0
Controller 2 Port 1 Status = Online

```

Bugzilla #449648

Red Hat discourages the use of **quota** on EXT3 file systems. This is because in some cases, doing so can cause a deadlock.

Testing has revealed that **kjournald** can sometimes block some EXT3-specific callouts that are used when **quota** is running. As such, Red Hat does not plan to fix this issue in Red Hat Enterprise Linux 4, as the modifications required would be too invasive.

Note that this issue is not present in Red Hat Enterprise Linux 5.

Bugzilla #451164

Hardware testing for the *Mellanox MT25204* has revealed that an internal error occurs under certain

high-load conditions. When the **ib_mthca** driver reports a catastrophic error on this hardware, it is usually related to an insufficient completion queue depth relative to the number of outstanding work requests generated by the user application.

Although the driver will reset the hardware and recover from such an event, all existing connections at the time of the error will be lost. This generally results in a segmentation fault in the user application. Further, if **opensm** is running at the time the error occurs, then you need to manually restart it in order to resume proper operation.

Bugzilla #452578

The Desktop Sharing connection icon displays its context menu when you double-click it, not when you right-click it. All other icons display their context menus when you right-click on them.

Bugzilla #451873

If the **ib_ehca** InfiniBand driver is loaded in port auto-detection mode (using module parameter **nr_ports=-1**), the IP-over-InfiniBand network interfaces (ibX) might become available too late. When this occurs, the **ifup ibX** command issued from the **openibd** startup script will fail; consequently, the ibX interface will not become available.

When this occurs, use the command **rcnetwork restart** to fix the problem.

Bugzilla #451873

In the *IBM Redbook "Implementing InfiniBand in IBM System p (SG247351) manual*, Table 6-3 (on page 220 of the PDF version) describes debug code bit definitions, where several HCA error indicator bits are also described.

Note that with eHCA2 adapters, bits 46 and 47 of these error indicator bits might return false positives.

Bugzilla #366961

On *HP ICH10* workstations, audio is only enabled through the front 3.5mm jacks. As such, to receive any audio output or use recording, you should plug in your headphones, speakers, or microphones to the front jacks. At present, the rear jacks, internal speaker, and master volume for this workstation do not work.

Bugzilla #429727

With this update, the default PCI detection and ordering mode for the following models have changed:

- *HP Proliant DL 580 G5*
- *HP Proliant DL 385 G2*
- *HP Proliant DL 585 G2*

These models use a device scanning and enumeration mode which is not the default for Red Hat Enterprise Linux 4 or 5. The mode used by these *HP Proliant* models could result in add-on cards being detected and added prior to onboard/internal devices. This unexpected ordering could cause difficulties when installing new instances of Red Hat Enterprise Linux, adding hardware, and maintenance.

The numbering of network interface cards (NIC) for the aforementioned *HP Proliant* models may change when they are updated with the Red Hat Enterprise Linux 4.7 kernel. The installer changes

NIC numbering if the **HWADDR=MAC ADDRESS** parameter is not defined in `/etc/sysconfig/network-scripts/ifcfg-eth[X]` for each installed NICs. As such, Red Hat recommends that you ensure this parameter is defined in order to avoid any problems arising from an unexpected NIC enumeration.

In addition, to avoid any NIC enumeration changes after updating these *HP Proliant* models to Red Hat Enterprise Linux 4.7, add the kernel boot parameter **pci=nobfsort** to `/boot/grub/grub.conf`.

Bugzilla #232499

When a volume group contains a mirror or snapshot, issuing the `lvchange` command with a volume group parameter may result in the following error messages:

```
Unable to change mirror log LV fail_secondary_mlog directly
Unable to change mirror image LV fail_secondary_mimage_0 directly
Unable to change mirror image LV fail_secondary_mimage_1 directly
```

These messages can be safely ignored.

Bugzilla #441870

Dell PowerEdge SC1435s systems may hang during boot-up. To avoid this, edit the **terminal** line in `grub.conf` and replace the string **serial console** with **console serial**.

Bugzilla #456533

The updated **ixgbe** driver does not support the *Intel 82598AT (Copper Pond 10GbE)*.

Bugzilla #454872

Red Hat Enterprise Linux 4.8 can detect online growing or shrinking of an underlying block device. However, there is no method to automatically detect that a device has changed size, so manual steps are required to recognize this and resize any file systems which reside on the given device(s). When a resized block device is detected, a message like the following will appear in the system logs:

```
VFS: busy inodes on changed media or resized disk sdi
```

If the block device was grown, then this message can be safely ignored. However, if the block device was shrunk without shrinking any data set on the block device first, the data residing on the device may be corrupted.

It is only possible to do an online resize of a filesystem that was created on the entire LUN (or block device). If there is a partition table on the block device, then the file system will have to be unmounted to update the partition table.

Bugzilla #479467

There is a known memory leak with the **res_n*** family of resolver routines (i.e. **res_nquery**, **res_nsearch** and **res_nmkquery**). Programs that use these functions will leak memory over time. It has been fixed in newer versions of `glibc`, however, the fix is too invasive to be applied to Red Hat Enterprise Linux 4. Programs that use these functions may need to be restarted occasionally to free memory.

Bugzilla #452513

The number of devices that can be handled during installation of Red Hat Enterprise Linux 4 depends on the size of the installation `initrd` image. Therefore, in situations where there are many devices attached to a machine (such as heavily populated Fibre Channel setups) installation will not

be possible unless number of visible devices is reduced.

Bugzilla #438895

The **aacraid** driver update that was first introduced in Red Hat Enterprise Linux 4.7 requires up to date **Adaptec PERC3/Di** firmware. Subsequent updates of Red Hat Enterprise Linux 4 (including this 4.8 update) require, that the **PERC3/Di** firmware is at version **2.8.1.7692, A13** or newer. The firmware may be obtained at the following location:

http://support.dell.com/support/downloads/download.aspx?c=us&cs=555&l=en&s=biz&releaseid=R168387&SystemID=PWE_PNT_PIII_1650&servicetag=&os=WNET1&impid=-1&formatcnt=4&libid=35&fileid=228550

Bugzilla #492371

During installation anaconda may not remove all the **Logical Volume Manager (LVM)** metadata that exists on a system prior to installation. This extra metadata may cause **LVM** tools to report missing volume groups or logical volumes after installation. To work around this issue, remove the stale **LVM** metadata after the installation is complete.

Bugzilla #481190

multipath does not silence the error messages printed by any of it's callout programs. Therefore, if **multipath** is run when paths are down, various error messages may be displayed. The messages that are displayed depend on the specific callout programs that **multipath** is using. For example, if **multipath** is run while there are failed scsi devices, **scsi_id** will print

```
lt;H>:<B>:<T>:<L>:Unable to get INQUIRY vpd 1 page 0x0.
lt;H>:<B>:<T>:<L>:sg_io failed status 0x0 0x1 0x0 0x0
```

Or, if **multipath -ll** is run while an *EMC CLARiiON* is down, the **mpath_prio_emc priority** callout will print **query command indicates error**

9.2. ia64 Architectures

Bugzilla #453033

On some *SGI Altix* systems that feature the IOC4 multi-function device, you may encounter problems when using attached IDE devices (such as CD-ROM drives). This is caused by a bug in the **sgiioc4** IDE driver, which prevents some devices from being detected properly on system boot.

You can work around this bug by manually loading the driver, which in turn allows attached IDE devices to be detected properly. To do so, run the following command as root:

```
/sbin/modprobe sgiioc4
```

A. REVISION HISTORY

Revision 2-4.400 Rebuild with publican 4.0.0	2013-10-31	Rüdiger Landmann
Revision 2-4 Rebuild for Publican 3.0	2012-07-18	Anthony Towns
Revision 2-3 BZ#627110 BZ#627111	Tue Feb 8 2011	Michael Hideo
Revision 1.0-0 Added Release Notes Updates for the General Availability (GA)	Thu May 07 2009	Ryan Lerch
Revision 0.1-0 Initial Version of the Release Notes	Thu May 07 2009	Ryan Lerch