



Red Hat Enterprise Linux 4

4.9 Release Notes

Release Notes for all Architectures

Red Hat Enterprise Linux 4 4.9 Release Notes

Release Notes for all Architectures

Ryan Lerch

Red Hat Engineering Content Services

rlersch@redhat.com

Legal Notice

Copyright © 2011 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

Release Notes for Red Hat Enterprise Linux 4.9

Table of Contents

1. INTRODUCTION	2
2. INSTALLATION	2
3. DRIVER UPDATES	3
4. GENERAL UPDATES	3
5. KERNEL-RELATED UPDATES	3
5.1. Bug Fixes	4
5.2. Enhancements	8
A. TECHNOLOGY PREVIEWS	9
B. KNOWN ISSUES	10
C. REVISION HISTORY	16

1. INTRODUCTION

Red Hat Enterprise Linux minor releases are an aggregation of individual enhancement, security and bug fix errata. The Red Hat Enterprise Linux 4.9 Release Notes documents the major changes made to the Red Hat Enterprise Linux 4 operating system and its accompanying applications for this minor release.

Detailed information on the changes that each updated package provides is available in the [Errata Documentation](#) for each package.

Red Hat Enterprise Linux 4 Lifecycle

Red Hat is releasing the 9th minor release update for Red Hat Enterprise Linux 4. This update transitions Red Hat Enterprise Linux to the Production III life cycle phase. This update is being released as an online only update on Red Hat Network. No new media kits are being issued. Red Hat Enterprise Linux 4 is now closed for new feature requests and new hardware enablement. Red Hat will continue to deliver security and bug fixes for Red Hat Enterprise Linux 4 until the End-of-Life scheduled for February 29, 2012. Customers requiring longer life cycle support should contact Red Hat or its support partners for Extended Life Cycle phase support programs.

2. INSTALLATION

Installation of Red Hat Enterprise Linux 4.9 is completed by upgrading from Red Hat Enterprise Linux 4.8.



IMPORTANT

Before upgrading, ensure that the system has the most recent version of the **up2date** tool installed by running the command **up2date up2date**



WARNING

Due to a known issue, if you are subscribed to the Red Hat Enterprise Linux 4 Extras channel and have the HelixPlayer package installed on your system, you must first manually remove the HelixPlayer package before updating the system to Red Hat Enterprise Linux 4.9. You can manually uninstall HelixPlayer by running the following command:

```
rpm -e HelixPlayer
```

Failing to uninstall the HelixPlayer package before updating your system may cause **up2date** to fail to update all packages to their Red Hat Enterprise Linux 4.9 versions.

To update a Red Hat Enterprise Linux 4.8 installation, first ensure that the system is subscribed to the 4.9 channel, then use the **up2date** command:

up2date --update

There is no media kit available for Red Hat Enterprise Linux 4.9.

3. DRIVER UPDATES

The **kernel-utils** package provides **microcode_ctl** utility code and the microcode data itself – supplied by Intel – to assist the kernel in updating the CPU microcode at system boot time. This microcode supports all current Intel x86- and Intel 64-based CPU models and takes advantage of the mechanism built-in to Linux that allows microcode to be updated after system boot. When loaded, the updated microcode corrects the behavior of various Intel processors, as described in processor specification updates issued by Intel for those processors. In the Red Hat Enterprise Linux 4.9 release the kernel-utils package has been updated with the 2010-09-14 version of Intel's microcode.

4. GENERAL UPDATES

Intel Xeon support in OProfile

OProfile is a system-wide Linux profiler, capable of running at low overhead. It consists of a kernel driver and a daemon for collecting raw sample data, along with a suite of tools for parsing that data into meaningful information. OProfile is generally used by developers to determine which sections of code consume the most amount of CPU time, and why.

OProfile in Red Hat Enterprise Linux 4.9 now supports Intel Xeon (formerly codenamed Nehalem) processor performance events.

Improved Checksum support in coreutils

Newer Red Hat disk images use the SHA-256 algorithm as the checksum method. Previously, verifying these checksums on Red Hat Enterprise Linux 4 was not a simple process. Red Hat Enterprise Linux 4.9 provides an updated **coreutils** package, containing checksum utilities for the whole SHA-2 family, allowing a user to easily count or verify SHA-224, SHA-256, SHA-384, and SHA-512 checksums.

KornShell

KSH-93 is the most recent version of the KornShell (**ksh**) by David Korn of AT&T Bell Laboratories, a shell programming language upwards-compatible with **sh** (the Bourne Shell). In Red Hat Enterprise Linux 4.9, **ksh** is updated to the upstream release version **ksh93t+**.

Firefox

Red Hat Enterprise Linux 4.9 features version 3.6.13 of the Mozilla Firefox open source web browser. Refer to the [Firefox 3.6 Release Notes](#) and [Firefox 3.6.13 Release Notes](#) for further details on the bug fixes and enhancements provided by this update.

SystemTap

SystemTap is a tracing and probing tool that allows users to study and monitor the activities of the operating system (particularly, the kernel) in fine detail. It provides information similar to the output of tools like netstat, ps, top, and iostat; however, SystemTap is designed to provide more filtering and analysis options for collected information.

SystemTap in Red Hat Enterprise Linux 4.9 has been updated to version 1.3. Most notably, this update provides support for the newer kprobe capabilities.

5. KERNEL-RELATED UPDATES

5.1. Bug Fixes

The updated Red Hat Enterprise Linux 4.9 kernel packages fix the following urgent- and high-priority bugs. This list of bug fixes is not comprehensive: refer to the "Bugs fixed" section of the RHN advisory for this kernel update for the complete list of bug fixes and enhancements.

BZ#586307

This change fixes forward time drift which has been observed with 64-bit Red Hat Enterprise Linux 4 virtual machines using PM timer based kernel tick accounting while running on KVM or HYPER-V hypervisor.

Virtual machines that are booted with the "divider=x" kernel parameter set to a value greater than 1 and that show the following line of text in the kernel boot messages are prone to the problem:

time.c: Using PM based timekeeping

However, this change also uncovered a bug in the Xen hypervisor, possibly causing backward time drift when this change is applied. If such time drift is observed with fully-virtualized Xen guests satisfying the aforementioned conditions, the kernel parameter **pmtimer_fine_grained=0** can be used as a temporary workaround until the hypervisor is updated.

BZ#537475

Red Hat Enterprise Linux 4 kernel 2.6.9-67 does not enable barriers if it detects WCE (Write Cache Enabled) as being disabled. However, the PERC 6/i controller on Dell 2950 hardware improperly reports WCE as being disabled while still maintaining a write cache. This could have caused file system corruption in the event of a power outage due to data being maintained in the controller cache even though the kernel was told that data was being written through to disk.

The solution to this issue provides proper software ordering, with the result that proper data integrity is achieved with devices that have no write caching (or for which write caching is disabled), and no command queuing. If command queuing or write caching is enabled, there is no guarantee of data integrity following a crash.

BZ#531192

Input entered on a PS/2 keyboard connected to certain Primergy servers, such as the Primergy TX120S1, was not registered. With this update, the PS/2 port is reset after loading PS/2 AUX, with the result that PS/2 keyboards work as expected.

BZ#494597

The Red Hat Enterprise Linux 4.8 kernel update contained a fix that had the side effect of causing the igb (Intel Gigabit Ethernet) driver to truncate some received packets when the Maximum Transmission Unit (MTU) was set to 2100 bytes or higher. The workaround for the issue was to set the network device's MTU to less than **2100**. This update fixes this flaw in the igb driver so that frames **2100** bytes or larger can once again be used.

BZ#504019

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.

BZ#554525

A regression prevented the Broadcom BCM5761 network device from working when in the first (top) PCI-E slot of Hewlett-Packard (HP) Z600 systems.

Note: The card worked in the 2nd or 3rd PCI-E slot.

BZ#574785

A race condition caused TX to stop in a guest using the virtio_net driver.

BZ#577378

A bug could have prevented NFSv3 clients from having the most up-to-date file attributes for files on a given NFSv3 file system. In cases where a file type changed, such as if a file was removed and replaced with a directory of the same name, the NFSv3 client may not have noticed this change until `stat(2)` was called (for example, by running `ls -l`).

BZ#594086

With multipath enabled, systems would occasionally halt when the `do_cciss_request` function was used. This was caused by wrongly-generated requests. Additional checks have been added to avoid the aforementioned issue.

BZ#616710

Erroneous pointer checks could have caused a kernel panic. This was due to a critical value not being copied when a network buffer was duplicated and consumed by multiple portions of the kernel's network stack. Fixing the copy operation resolved this bug.

BZ#619449

In some situations a bug prevented "force online" succeeding for a DASD device.

BZ#619527

Previously, allocating fallback cqr for DASD reserve/release IOCTLS failed because it used the memory pool of the respective device. This update preallocates sufficient memory for a single reserve/release request.

BZ#528066

A bug in the `mptctl_do_mpt_command()` function in the mpt driver may have resulted in crashes during boot on i386 systems with certain adapters using the mpt driver, and also running the hugemem kernel.

BZ#529037

On certain hardware, the igb driver was unable to detect link statuses correctly. This may have caused problems for network bonding, such as failover not occurring.

BZ#529063

A fix in the Red Hat Security Advisory [RHSA-2009:1024](#) introduced a regression which caused the `diskdump` command to fail on systems with certain adapters using the `qla2xxx` driver. With this update, `diskdump` now works as expected.

BZ#532593

A fix in the Red Hat Security Advisory [RHSA-2009:1024](#) introduced a regression. After updating to

Red Hat Enterprise Linux 4.8 and rebooting, network links often failed to be brought up for interfaces using the forcedeth driver. Additionally, **no link during initialization** messages may have been logged.

BZ#539506

If a process was using **ptrace()** to trace a multi-threaded process, and that multi-threaded process dumped its core, the process performing the trace could hang in **wait4()**. This issue could be triggered by running **strace -f** on a multi-threaded process that was dumping its core, resulting in the strace command hanging.

BZ#523691

Using the bnx2 driver under heavy NFS usage could have caused the kernel to panic; this was traced back to the **poll_freewait()** function. The following message was received prior to the kernel panic:

RPC: Invalid TCP record fragment length

This update resolves this issue by updating the 5706/5708 firmware, eliminating TSO header modifications, and fixing jumbo frame error handling, with the result that the bnx2 driver no longer has to modify TCP/IP header fields when transmitting TCP Segmentation Offload (TSO) packets.

BZ#544381

In some circumstances, when a Red Hat Enterprise Linux client connected to a re-booted Windows-based NFS server, server-side filehandle-to-inode mapping changes caused a kernel panic.

bad_inode_ops handling was changed to prevent this.

Note: filehandle-to-inode mapping changes may still cause errors, but not panics.

BZ#547213

Under some circumstances, a locking bug could have caused an online ext3 file system resize to deadlock, which may have, in turn, caused the file system or the entire system to become unresponsive. In either case, a reboot was required after the deadlock. With this update, using **resize2fs** to perform an online resize of an ext3 file system works as expected.

BZ#477784

On 32-bit architectures, if a file was held open and frequently written for more than **25** days, it was possible that the kernel would stop flushing those writes to storage.

BZ#483783

Attaching a certain USB device could have caused a kernel crash due to a divide-by-zero error.

BZ#491147

A bug in the **timer_interrupt()** function may have caused the system time to move up to two days or more into the future, or to be delayed for several minutes. This bug only affected Intel 64 and AMD64 systems that have the High Precision Event Timer (HPET) enabled in the BIOS, and could have caused problems for applications that require timing to be accurate.

BZ#494404

Due to insufficient memory barriers in the network code, a process sleeping in the **select()** function may have missed notifications about new data. In rare cases, this bug may have caused a process to sleep forever.

BZ#494428

The **tcp_ack()** function cleared the **probes_out** variable even if there were outstanding packets. When low TCP keepalive intervals were used, this bug may have caused problems, such as connections terminating, when using remote tools such as rsh and rlogin.

BZ#494688

Using **netdump** may have caused a kernel deadlock on some systems.

BZ#495930

The driver version number in the ata_piix driver was not changed between Red Hat Enterprise Linux 4.7 and Red Hat Enterprise Linux 4.8, even though changes had been made between these releases. This could have prevented the driver from loading on systems that check driver versions, as this driver appeared older than it was. This update increases the driver version number.

BZ#496120

An error in the MPT Fusion driver makefile caused CSMI ioctls to not work with Serial Attached SCSI devices.

BZ#496201

Off-by-one errors in the time normalization code could have caused the **clock_gettime()** function to return one billion nanoseconds, rather than adding an extra second. This bug could have caused the name service cache daemon (**nscd**) to consume excessive CPU resources.

BZ#504080

Time-outs resulted in I/O errors being logged to **/var/log/messages** when running **mt erase** on tape drives using certain LSI MegaRAID SAS adapters, preventing the command from completing. The megaraid_sas driver's timeout value is now set to the OS layer value.

BZ#505122

A bug in the SCSI implementation caused **Aborted Command - internal target failure** errors to be sent to Device-Mapper Multipath, without retries, resulting in Device-Mapper Multipath marking the path as failed and making a path group switch. With this update, all errors that return a sense key in the SCSI mid layer (including **Aborted Command - internal target failure**) are retried.

BZ#507680

The Emulex LPFC driver has been updated to version 8.0.16.47, which fixes a memory leak that caused memory allocation failures and system hangs.

BZ#515258

On 64-bit PowerPC systems, a rollover bug in the ibmveth driver could have caused a kernel panic. In a reported case, this panic occurred on a system with a large uptime and under heavy network load.

BZ#517329

The data buffer that the **ethtool_get_strings()** function allocated, for the igb driver, was smaller than the amount of data that was copied in **igb_get_strings()**, because of a miscalculation in **IGB_QUEUE_STATS_LEN**, resulting in memory corruption. This bug could have led to a kernel panic.

BZ#507680

The Emulex LPFC driver has been upgraded to an updated upstream version, which fixes a memory leak that caused memory allocation failures and system hangs.

BZ#503489

A locking issue caused the qla2xxx ioctl module to hang after encountering errors. This locking issue has been corrected. This ioctl module is used by the QLogic SAN management tools, such as SANsurfer and scli.

BZ#548496

It was found that the **lpfc_find_target()** function could loop continuously when scanning a list of nodes due to a missing spinlock. This missing spinlock allowed the list to be changed after the **list_empty()** test, resulting in a **NULL** value, causing the loop. This update adds the spinlock, resolving the issue.

BZ#577178

A bug was found in the way the megaraid_sas driver handled physical disks and management IOCTLS. All physical disks were exported to the disk layer, allowing an oops in the **megasas_complete_cmd_dpc()** function when completing the IOCTL command if a timeout occurred.

5.2. Enhancements

The updated Red Hat Enterprise Linux 4.9 kernel packages also provide the following enhancement:

BZ#553745

Support for the Intel architectural performance monitoring subsystem (arch_perfmon). On supported CPUs, arch_perfmon offers means to mark performance events and options for configuring and counting these events.

A. TECHNOLOGY PREVIEWS

Technology Preview features are currently *not* supported under Red Hat Enterprise Linux 4.9 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.

B. KNOWN ISSUES



NOTE

The *4.9 Release Notes* were updated on April 22, 2011 with the first three descriptions in the list below: BZ#588753, BZ#509095 and BZ#487443. These issues are part of the [RHBA-2011:0243 device-mapper-multipath bug fix and enhancement update](#) .

BZ#588753

Installing the device-mapper-multipath RPM package could have caused the default installed `/etc/kpartx.conf` configuration file to be modified. As a consequence, verifying the package with the `rpm -V` command failed due to the `/etc/kpartx.conf` file's size, md5sum, and last modification time having been changed. This update adds flags to the device-mapper-multipath package's spec file that inform rpm that the size, md5sum and modification time of the `/etc/kpartx.conf` file may change, with the result that verifying the package now succeeds in this situation.

BZ#509095

When the system boots, multipath reads the WWID (World Wide Identifier) of the storage device on each path, and creates a binding from each WWID to an `mpath[n]` device name. By default, these bindings are written to the `/var/lib/multipath/bindings` file, which is referenced in all future scans to ensure that the binding from an `mpath[n]` to a storage device WWID remains constant for the life of the system. When the `/var` directory was mounted on a separate file system from root ("`/`"), then `/var` could have been mounted after multipath had created the initial bindings file. In this situation, the next time multipath ran, it did not find a bindings file and therefore proceeded to create new bindings, which it wrote to the `/var/lib/multipath/bindings` file.

However, if the scan order differed, then the bindings would change, thus causing one or more `multipath[n]` names to point to different storage devices. This could have resulted in: inconsistent multipath maps; multiple maps with the same WWID; maps with the wrong paths; and potentially data loss. Note that the possibility that the scan order can potentially differ is greater in an environment where LUNs (Logical Unit Numbers) are added or removed after the system boots.

The `multipath.conf` configuration file now includes a default `bindings_file` directive. On installations on which the `/var/lib/multipath/bindings` file is not located on the same device as the root file system, then this directive should specify a custom path to a bindings file on the same device as the root file system, for multipath to use instead of the default. By setting this directive, multipath will use the same bindings during boot as it does during normal operation, multipath devices will remain consistent, and the possibility of data loss is avoided.

For more information on the `bindings_file` directive, refer to the `/usr/share/doc/device-mapper-multipath-0.4.5/multipath.conf.annotated` file.

BZ#487443

Multipath is unable to reconfigure a multipath device while the `multipathd` service is running, or add new paths to the device, because the kernel does not allow multipath to reload a table with a failed path in it. In order to remove a failed path from the device table, the following procedure should be followed as a workaround:

1. Stop the `multipathd` service:

```
# service multipathd stop
```

2. Run the `multipath` command in order to reconfigure the device.

3. Start multipathd:

```
# service multipathd start
```

Running the **multipath** command while the multipathd service is stopped causes multipath to remove the failed path from the device table when it reconfigures, thus allowing the new path to be added. However, when the failed path comes back online, it will no longer be part of the multipath device. The **multipath** command must be manually run again in order to reintegrate this device. Note that if there are no failed paths, it is not necessary to stop the **multipathd** service before running multipath.

BZ# 614559

When using the sky2 driver for Marvell Ethernet adapters with the following device:

Table B.1. Device Identifier

Ethernet controller	Marvell Technology Group Ltd.
	88E8052 PCI-E ASF Gigabit Ethernet Controller (rev 20)
Vendor ID	11ab
Device ID	4360

...packet transmission may time out with the following message being written to netconsole: **sky2 eth1: tx timeout.**

This is a known issue: after receiving these messages, the system must be rebooted in order to fix the packet transmission issues.

BZ#523255

When dm-multipath is used on a storage device that implements ALUA, and **group-by-prio** is enabled, then the path groups are established when the device is configured. The paths with the same priority are grouped together, the group priority is calculated as the sum of the path priorities, and the path group with the highest sum is selected for I/O. If a path's priority changes, the group priority is re-calculated, and the active path group may change. The path grouping is not changed, even though some members of the group may now have different priorities. If you wish to re-establish the path grouping after a change, then you must enter the command

```
multipathd -k"reconfigure"
```

BZ#653424

When dm-multipath is used on a storage device that implements ALUA, and **group-by-prio** is enabled, then the path groups are established when the device is configured. The paths with the same priority are grouped together, the group priority is calculated as the sum of the path priorities, and the path group with the highest sum is selected for I/O. If a path's priority changes, the group priority is re-calculated, and the active path group may change. The path grouping is not changed, even though some members of the group may now have different priorities. If you wish to re-establish the path grouping after a change, then you must enter the command

```
multipathd -k"reconfigure"
```

BZ#653424

Broken dependencies in the `up2date-4.8.1-33.el4_8.9` package may result in an update of the `up2date` package failing. Consequently, if the `up2date-4.8.1-33.el4_8.9` package is installed on a system the new `up2date` package must be manually downloaded from the Red Hat Network and installed using the `rpm` command.

BZ#459785

Japanese language JP106 keyboards will not function correctly when booting into Rescue Mode on **Red Hat Enterprise Linux 4.9**.

BZ#494022

Updating all packages from **Red Hat Enterprise Linux 4.8** to **Red Hat Enterprise Linux 4.9** 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
```

BZ#443137

Some x86 64-bit systems may hang during boot when the `noapic` debug kernel parameter is used.

BZ#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.

BZ#484117

The Logical Volume Manager in Red Hat Enterprise Linux 4.9 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.

BZ#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.

BZ#468097

On systems where the BIOS is able to do both legacy (`acpiphp`) and native (`pciehpc`) 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**.

BZ#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.

BZ#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.

BZ#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.

BZ#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.

BZ#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.

BZ#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`.

BZ#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.

BZ#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**.

BZ#456533

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

BZ#454872

Red Hat Enterprise Linux 4.9 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.

BZ#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.

BZ#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.

BZ#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.

BZ#481190

multipath does not silence the error messages printed by any of its 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**

BZ#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
```

C. REVISION HISTORY

Revision 0-15.400 Rebuild with publican 4.0.0	2013-10-31	Rüdiger Landmann
Revision 0-15 Rebuild for Publican 3.0	2012-07-18	Anthony Towns
Revision 1.0-0 Added Release Notes Updates for the General Availability (GA)	Tue Dec 07 2010	Ryan Lerch