
Red Hat Enterprise Linux 5.4 Release Notes

Release Notes for all architectures.

Copyright © 2009 Red Hat, Inc.. This material may only be distributed subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version of the OPL is presently available at <http://www.opencontent.org/openpub/>).

Red Hat and the Red Hat "Shadow Man" logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

All other trademarks referenced herein 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
PO Box 13588 Research Triangle Park, NC 27709 USA

Abstract

This document details the Release Notes for Red Hat Enterprise Linux 5.4.

1. Virtualization Updates	2
2. Clustering Updates	3
2.1. Fencing Improvements	4
3. Networking Updates	5
4. Filesystems and Storage updates	5
5. Desktop Updates	6
5.1. Advanced Linux Sound Architecture	6
5.2. Graphics Drivers	6
5.3. Laptop Support	7
6. Tools Updates	7
7. Architecture Specific Support	8
7.1. i386	8
7.2. x86_64	9
7.3. PPC	9
7.4. s390	9
8. Kernel Updates	11
8.1. General Kernel Feature Support	11
8.2. General Platform Support	13

8.3. Driver Updates 13
9. Technology Previews 18
A. Revision History 19

This document contains the Release Notes for the Red Hat Enterprise Linux 5.4 family of products including:

- Red Hat Enterprise Linux 5 Advanced Platform for x86, AMD64/Intel® 64, Itanium Processor Family, System p and System z
- Red Hat Enterprise Linux 5 Server for x86, AMD64/Intel® 64, Itanium Processor Family, System p and System z
- Red Hat Enterprise Linux 5 Desktop for x86 and AMD64/Intel®

The Release Notes provide high level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 5.4. For detailed documentation on all changes to Red Hat Enterprise Linux for the 5.4 update, refer to the [Technical Notes](#)¹

1. Virtualization Updates

Red Hat Enterprise Linux 5.4 now includes full support for the Kernel-based Virtual Machine (KVM) hypervisor on x86_64 based architectures. KVM is integrated into the Linux kernel, providing a virtualization platform that takes advantage of the stability, features, and hardware support inherent in Red Hat Enterprise Linux. Virtualization using the KVM hypervisor is supported on wide variety of guest operating systems, including:

- Red Hat Enterprise Linux 3
- Red Hat Enterprise Linux 4
- Red Hat Enterprise Linux 5
- Windows XP
- Windows Server 2003
- Windows Server 2008

 **Important**
Xen based virtualization is fully supported. However, Xen-based virtualization requires a different version of the kernel to function. The KVM hypervisor can only be used with the regular (non-Xen) kernel.

 **Warning**
While Xen and KVM may be installed on the same system, the default networking configuration for these are different. Users are strongly recommended to only install one hypervisor on a system.

¹ http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.4/html/Technical_Notes/



Note

Xen is the default hypervisor that is shipped with Red Hat Enterprise Linux. As such all configuration defaults are tailored for use with the Xen hypervisor. For details on configuring a system for KVM, please refer to the Virtualization Guide.

Virtualization using KVM allows both 32-bit and 64-bit versions of guest operating systems to be run without modification. Paravirtualized disk and network drivers have also been included in Red Hat Enterprise Linux 5.4 for enhanced I/O performance. All the libvirt based tools (i.e. **virsh**, **virt-install** and **virt-manager**) have also been updated with added support for KVM.

USB passthrough with the KVM hypervisor is considered to be a Technology Preview for the 5.4 release.

With resolution of various issues such as: save/restore, live migration and core dumps, Xen based 32 bit paravirtualized guests on x86_64 hosts are no longer classed as a Technology Preview, and are fully supported on Red Hat Enterprise Linux 5.4.

the **etherboot** package has been added in this update, providing the capability to boot guest virtual machines using the Preboot eXecution Environment (PXE). This process occurs before the OS is loaded and sometimes the OS has no knowledge that it was booted through PXE. Support for etherboot is limited to usage in the KVM context.

The **qspice** packages have been added to Red Hat Enterprise Linux 5.4 to support the *spice protocol* in **qemu-kvm** based virtual machines. **qspice** contains both client, server and web browser plugin components. However, only the **qspice** server in the **qspice-libs** package is fully supported. The qspice client (supplied by the qspice package) and qspice mozilla plugin (supplied by the qspice-mozilla package) are both included as Technology Previews. The **qspice-libs** package contains the server implementation that is used in conjunction with **qemu-kvm** and as such is fully supported. However, in Red Hat Enterprise Linux 5.4 there is no **libvirt** support for the spice protocol; the only supported use of **spice** in Red Hat Enterprise Linux 5.4 is through the use of the Red Hat Enterprise Virtualization product.



Important

The virtio-win component is only available via the Red Hat Network, and is *not* included on the physical Supplementary CD for Red Hat Enterprise Linux 5.4. For more information, see the *Red Hat Knowledgebase*².

2. Clustering Updates

Clusters are multiple computers (nodes) working in concert to increase reliability, scalability, and availability to critical production services.

All updates to clustering in Red Hat Enterprise Linux 5.4 are detailed in the Technical Notes. Further information on clustering in Red Hat Enterprise Linux is available in the *Cluster Suite Overview*³ and the *Cluster Administration*⁴ documents.

³ http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Cluster_Suite_Overview/index.html

⁴ http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Cluster_Administration/index.html

Cluster Suite tools have been upgraded to support automatic hypervisor detection. However, running the cluster suite in conjunction with KVM hypervisor is considered to be a Technology Preview.

OpenAIS now provides broadcast network communication in addition to multicast. This functionality is considered Technology Preview for standalone usage of OpenAIS and for usage with the Cluster Suite. Note, however, that the functionality for configuring OpenAIS to use broadcast is not integrated into the cluster management tools and must be configured manually.



Note

SELinux in Enforcing mode is not supported with the Cluster Suite; Permissive or Disabled modes must be used. Using Cluster Suite on bare metal PPC systems is not supported. Guests running Cluster Suite on VMWare ESX hosts and using `fence_vmware` is considered a Technology Preview. Running Cluster Suite in guests on VMWare ESX hosts that are managed by Virtual Center is not supported.

Mixed architecture clusters using Cluster Suite are not supported. All Nodes in the cluster must be of the same architecture. For the purposes of Cluster Suite, x86_64, x86 and ia64 are considered to be the same architecture, so running clusters with combinations of these architectures is supported.

2.1. Fencing Improvements

Fencing is the disconnection of a node from the cluster's shared storage. Fencing cuts off I/O from shared storage, thus ensuring data integrity.

In Red Hat Enterprise Linux 5.4, fencing support on Power Systems has been added, as a Technology Preview, for IBM Logical Partition (LPAR) instances that are managed using the Hardware Management Console (HMC) ([BZ#485700](#))⁵. Fencing support has also been added, as a Technology Preview for Cisco MDS 9124 & Cisco MDS 9134 Multilayer Fabric Switches ([BZ#480836](#)).⁶

The `fence_virsh` fence agent is provided in this release of Red Hat Enterprise Linux as a Technology Preview. `fence_virsh` provides the ability for one guest (running as a domU) to fence another using the libvirt protocol. However, as `fence_virsh` is not integrated with cluster-suite it is not supported as a fence agent in that environment.

The `fence_scsi` man page has been updated, detailing the following limitations:

The `fence_scsi` fencing agent requires a minimum of three nodes in the cluster to operate. For an FC connected SAN devices, these must be physical nodes. SAN devices connected via iSCSI may use virtual or physical nodes. In addition, `fence_scsi` cannot be used in conjunction with `qdisk`.

Additionally, the following new articles on fencing have been published on the Red Hat Knowledge Base:

- SCSI Fencing (Persistent Reservations) with Red Hat Enterprise Linux 5 Advanced Platform Cluster Suite: <http://kbase.redhat.com/faq/docs/DOC-17809>
- Using `fence_vmware` with Red Hat Enterprise Linux 5 Advanced Platform Cluster Suite: <http://kbase.redhat.com/faq/docs/DOC-17345>

⁵ https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=485700

⁶ https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=480836

3. Networking Updates

With this update, Generic Receive Offload (GRO) support has been implemented in both the kernel and the userspace application, **ethtool**.^{([BZ#499347](https://bugzilla.redhat.com/show_bug.cgi?id=499347))}⁷ The GRO system increases the performance of inbound network connections by reducing the amount of processing done by the Central Processing Unit (CPU). GRO implements the same technique as the Large Receive Offload (LRO) system, but can be applied to a wider range of transport layer protocols. GRO support has also been added to a several network device drivers, including the `igb` driver for Intel® Gigabit Ethernet Adapters and the `ixgbe` driver for Intel 10 Gigabit PCI Express network devices.

The Netfilter framework (the portion of the kernel responsible for network packet filtering) has been updated with added support for Differentiated Services Code Point (DSCP) values

the **bind** (Berkeley Internet Name Domain) package provides an implementation of the DNS (Domain Name System) protocols. Previously, `bind` did not offer a mechanism to easily distinguish between requests that will receive authoritative and non-authoritative replies. Consequently, an incorrectly configured server may have replied to requests that should have been denied. With this update, `bind` has been updated, providing the new option **allow-query-cache** that controls access to non-authoritative data on a server (for example: cached recursive results and root zone hits).^{([BZ#483708](https://bugzilla.redhat.com/show_bug.cgi?id=483708))}⁸

4. Filesystems and Storage updates

In the 5.4 update, several significant additions have been made to file systems support. Base Red Hat Enterprise Linux now includes the **Filesystem in Userspace (FUSE)** kernel modules and user space utilities, allowing users to install and run their own **FUSE** file systems on an unmodified Red Hat Enterprise Linux kernel ([BZ#457975](https://bugzilla.redhat.com/show_bug.cgi?id=457975))⁹. Support for the **XFS** file system has also been added to the kernel for future product enablement ([BZ#470845](https://bugzilla.redhat.com/show_bug.cgi?id=470845))¹⁰. The FIEMAP input/output control (`ioctl`) interface has been implemented, allowing the physical layout of files to be mapped efficiently. The FIEMAP `ioctl` can be used by applications to check for fragmentation of a specific file or to create an optimized copy of a sparsely allocated file ([BZ#296951](https://bugzilla.redhat.com/show_bug.cgi?id=296951))¹¹.

Additionally, the Common Internet File System (CIFS) has been updated in the kernel ([BZ#465143](https://bugzilla.redhat.com/show_bug.cgi?id=465143))¹². The ext4 file system (included in Red Hat Enterprise Linux as a Technology Preview) has also been updated ([BZ#485315](https://bugzilla.redhat.com/show_bug.cgi?id=485315))¹³.

In Red Hat Enterprise Linux 5.4, the use of the Global File System 2 (GFS2) as a single server file system (i.e. not in a clustered environment) is deprecated. Users of GFS2 that do not need high availability clustering are encouraged to look at migrating to other file systems like the ext3 or xfs offerings. The xfs file system is specifically targeted at very large file systems (16 TB and above). Existing users will continue to be supported.

The required semantics indicate that a process which completes a **stat**, **write**, **stat**, should see a different **mtime** (time of last modification) on the file in the results from the second `stat` call compared to the **mtime** in the results from the first `stat` call. File times in NFS are maintained strictly by the server, so the file **mtime** will not be updated until the data has been transmitted to the server

⁷ https://bugzilla.redhat.com/show_bug.cgi?id=499347

⁸ https://bugzilla.redhat.com/show_bug.cgi?id=483708

⁹ https://bugzilla.redhat.com/show_bug.cgi?id=457975

¹⁰ https://bugzilla.redhat.com/show_bug.cgi?id=470845

¹¹ https://bugzilla.redhat.com/show_bug.cgi?id=296951

¹² https://bugzilla.redhat.com/show_bug.cgi?id=465143

¹³ https://bugzilla.redhat.com/show_bug.cgi?id=485315

via the **WRITE NFS** protocol operation. Simply copying data into the pagecache is not sufficient to cause the **mtime** to be updated. This is one place where NFS differs from local file systems. Therefore, an NFS filesystem which is under a heavy write workload may result in stat calls having a high latency.[\(BZ#469848\)](#)¹⁴

The ext4 filesystem Technology Preview has been refreshed with updated userspace tools. Ext4 is an incremental improvement on the ext3 file system developed by Red Hat and the Linux community.



Note

In previous versions of Red Hat Enterprise Linux utilizing the ext4 Technology Preview, ext4 filesystems were labeled as **ext4dev**. With this update, ext4 filesystems are now tagged as **ext4**.

With this update, the **dmraid** logwatch-based email reporting feature has been moved from the **dmraid-events** package into the new **dmraid-events-logwatch** package. Consequently, systems that use this dmraid feature will need to complete the following manual procedure:

1. ensure the new 'dmraid-events-logwatch' package is installed.
2. un-comment the functional portion of the `/etc/cron.d/dmeventd-logwatch` crontab file.

[\(BZ#512833\)](#)¹⁵

samba3x and ctdb are provided as a Technology Preview on the x86_64 platform. Samba3x package provides Samba 3.3 and ctdb provide a clustered TDB backend. Running samba3x and ctdb on a set of cluster nodes with GFS file system will allow the export of a clustered CIFS filesystem.



Important

the samba3x packages conflict with the samba-3.0 packages shipped with Red Hat Enterprise Linux 5. To use the Technology Preview, it is recommended to perform a fresh install that does not include the samba-3.0 packages, and then to install the samba3x packages from the Supplementary media.

5. Desktop Updates

5.1. Advanced Linux Sound Architecture

In Red Hat Enterprise Linux 5.4, the Advanced Linux Sound Architecture (ALSA) has been updated — providing enhanced support for High Definition Audio (HDA).

5.2. Graphics Drivers

The **ati** driver for ATI video devices has been updated.

The **i810** and **intel** drivers for Intel integrated display devices have been updated.

The **mga** driver for Matrox video devices has been updated.

¹⁴ https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=469848

¹⁵ https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=512833

The **nv** driver for nVidia video devices has been updated.

5.3. Laptop Support

Previously, when undocking and docking some laptops with docking stations containing integrated CD/DVD drives, the drive would no longer be recognized. The system would need to be rebooted for the drive to be accessible. With this update, the ACPI docking drivers have been updated in the kernel, resolving this issue. ([BZ#485181](#)).¹⁶

6. Tools Updates



Important

All the IBM Java components are available online due to a late detection of missing COPYRIGHT notice. This applies to the Supplementary CD contents for Red Hat Enterprise Linux 5 on all architectures and releases. For more information, see the *Red Hat Knowledgebase*¹⁷.

- **SystemTap** is now fully supported, and has been re-based to the latest upstream version. This update features improved user-space probing through shared libraries, experimental DWARF unwinding, and a new `<sys/sdt.h>` header file which provides dtrace-compatible markers.

This re-base also enhances support for **debuginfo-less** operations. Typecasting (through the `@cast` operator) is now supported, along with kernel tracepoint probing. Several '**kprobe.***' probe bugs that hampered **debuginfo-less** operations are also now resolved.

SystemTap also features several documentation improvements. A new '**3stap**' feature provides users with useful man pages on most SystemTap probes and functions. The **systemtap-testsuite** package also features a larger library of sample scripts.

For more information about the SystemTap re-base, please refer to the SystemTap section Package Updates chapter of the Technical Notes.

- Systemtap tracepoints are placed in important sections of the kernel, allowing system administrators to analyze the performance of, and debug portions of code. In Red Hat Enterprise Linux 5.4, tracepoints have been added to the following sections of the kernel subsystem as a Technology Preview:
 - Memory Management (mm) ([BZ#493444](#))¹⁸
 - Block Device I/O (blktrace)([Bugzilla #493454](#))¹⁹
 - Network File System (NFS) ([BZ#499008](#))²⁰
 - Page Cache and Networking stacks ([BZ#475719](#))²¹
 - Scheduler ([BZ#497414](#))²²
- The Gnu Compiler Collection version 4.4 (GCC4.4) is now included in this release as a Technology Preview. This collection of compilers include C, C++, and Fortran compilers along with support libraries.

¹⁶ https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=485181

- **glibc new MALLOC behaviour:** The upstream glibc has been changed recently to enable higher scalability across many sockets and cores. This is done by assigning threads their own memory pools and by avoiding locking in some situations. The amount of additional memory used for the memory pools (if any) can be controlled using the environment variables `MALLOC_ARENA_TEST` and `MALLOC_ARENA_MAX`.

`MALLOC_ARENA_TEST` specifies that a test for the number of cores is performed once the number of memory pools reaches this value. `MALLOC_ARENA_MAX` sets the maximum number of memory pools used, regardless of the number of cores.

The glibc in the RHEL 5.4 release has this functionality integrated as a Technology Preview of the upstream malloc. To enable the per-thread memory pools the environment variable `MALLOC_PER_THREAD` needs to be set in the environment. This environment variable will become obsolete when this new malloc behaviour becomes default in future releases. Users experiencing contention for the malloc resources could try enabling this option.

7. Architecture Specific Support

7.1. i386

- In a virtual environment, timekeeping for Red Hat Enterprise Linux 64-bit kernels can be problematic, since time is kept by counting timer interrupts. De- and re-scheduling the virtual machine can cause a delay in these interrupts, resulting in a timekeeping discrepancy. This kernel release reconfigures the timekeeping algorithm to keep time based on a time-elapsed counter. ([Bugzilla #463573](#)²³)
- It was found that, if their stacks exceed the combined size of ~4GB, 64-bit threaded applications slowed down drastically in `pthread_create()`. This is because **glibc** uses `MAP_32BIT` to allocate those stacks. As the use of `MAP_32BIT` is a legacy implementation, this update adds a new flag (`MAP_STACK mmap`) to the kernel to avoid constraining 64-bit applications. ([Bugzilla #459321](#)²⁴)
- The update includes a feature bit that encourages TSCs to keep running in deep-C states. This bit `NONSTOP_TSC` acts in conjunction with `CONSTANT_TSC`. `CONSTANT_TSC` indicates that the TSC runs at constant frequency irrespective of P/T- states, and `NONSTOP_TSC` indicates that TSC does not stop in deep C-states. ([Bugzilla #474091](#)²⁵)
- This update includes a patch to include `asm-x86_64` headers in kernel-devel packages built on or for i386, i486, i586 and i686 architectures. ([Bugzilla #491775](#)²⁶)
- This update includes a fix to ensure that specifying `memmap=X$Y` as a boot parameter on i386 architectures yields a new BIOS map. ([Bugzilla #464500](#)²⁷)
- This update adds a patch to correct a problem with the Non-Maskable Interrupt (NMI) that appeared in previous kernel releases. The problem appeared to affect various Intel processors and caused the system to report the NMI watchdog was 'stuck'. New parameters in the NMI code correct this issue. ([Bugzilla #500892](#)²⁸)
- This release re-introduces PCI Domain support for HP xw9400 and xw9300 systems. ([Bugzilla #474891](#)²⁹)
- Functionality has been corrected to export module `powernow-k8` parameters to `/sys/modules`. This information was previously not exported. ([Bugzilla #492010](#)³⁰)

7.2. x86_64

- An optimization error was found in **linux-2.6-misc-utrace-update.patch**. When running 32-bit processes on a 64-bit machine systems didn't return ENOSYS on missing (out of table range) system calls. This kernel release includes a patch to correct this. ([Bugzilla #481682](#)³¹)
- Some cluster systems were found to boot with an unstable time source. It was determined that this was a result of kernel code not checking for a free performance counter (*PERFCTR*) when calibrating the *TSC* (Time Stamp Clock) during the boot process. This resulted, in a small percentage of cases, in the system defaulting to a busy *PERFCTR* and getting unreliable calibrations.

A fix was implemented to correct this by ensuring the system checked for a free *PERFCTR* before defaulting ([Bugzilla #467782](#)³²). This fix, however, cannot satisfy all possible contingencies as it is possible that all *PERFCTR*s will be busy when required for *TSC* calibration. Another patch has been included to initiate a kernel panic in the unlikely event (fewer than 1% of cases) that this scenario occurs. ([Bugzilla #472523](#)³³).

7.3. PPC

- This kernel release includes various patches to update the *spufs* (Synergistic Processing Units file system) for Cell processors. ([Bugzilla #475620](#)³⁴)
- An issue was identified wherein `/proc/cpuinfo` would list logical PVR Power7 processor architecture as "unknown" when `show_cpuinfo()` was run. This update adds a patch to have `show_cpuinfo()` identify Power7 architectures as Power6. ([Bugzilla #486649](#)³⁵)
- This update includes several patches that are required to add/improve MSI-X (Message Signaled Interrupts) support on machines using System P processors. ([Bugzilla #492580](#)³⁶)
- A patch has been added to this release to enable the functionality of the previously problematic power button on Cell Blades machines. ([Bugzilla #475658](#)³⁷)

7.4. s390

Red Hat Enterprise Linux introduces a wide range of new features for IBM System z machines, most notably:

- Utilizing Named Saved Segments (NSS), the z/VM hypervisor makes operating system code in shared real memory pages available to z/VM guest virtual machines. With this update, multiple Red Hat Enterprise Linux guest operating systems on the z/VM can boot from the NSS and be run from a single copy of the Linux kernel in memory. ([BZ#474646](#))³⁸
- Device driver support has been added in this update for the new IBM System z PCI cryptography accelerators, utilizing the same interfaces as prior versions. ([BZ#488496](#))³⁹
- Red Hat Enterprise Linux 5.4 adds support for processor degradation, which allows processor speed to be reduced in some circumstances (i.e. system overheating). ([BZ#474664](#))⁴⁰ This new feature allows automation software to observe the machine state and act based on defined policies.



Note

Processor degradation is supported on z990, z890 and later systems and is observed through SCLP system service event type 4 event qualifier 3. STSI will report the

new capacity of the processor in the file: `/sys/devices/system/cpu/cpuN/capability`.

- Control Program Identification (CPI) descriptive data is used to identify individual systems on the Hardware Management Console (HMC). With this update, CPI data can now be associated with a Red Hat Enterprise Linux instance. (BZ#475820)⁴¹

For more information on CPI refer to the [Device Drivers, Features, and Commands document](#)⁴²

- Fibre Channel Protocol (FCP) performance data can now be measured on Red Hat Enterprise Linux instances on the IBM System z platform. (BZ#475334)⁴³ Metrics that are collected and reported on include:
 - Performance relevant data on stack components such as Linux devices, Small Computer System Interface (SCSI) Logical Unit Numbers (LUNs) and Host Bus Adapter (HBA) storage controller information.
 - Per stack component: current values of relevant measurements as throughput, utilization and other applicable measurements.
 - Statistical aggregations (minimum, maximum, averages and histogram) of data associated with I/O requests including size, latency per component and totals.
- Support has been added to the kernel to issue EMC Symmetrix Control I/O. This update provides the ability to manage EMC Symmetrix storage arrays with Red Hat Enterprise Linux on the IBM System z platform. (BZ#461288)⁴⁴
- A new feature has been implemented in the kernel to perform an Initial Program Load (IPL) on a Red Hat Enterprise Linux virtual machine immediately following a kernel panic and dump. (BZ#474688)⁴⁵
- Hardware that supports the configuration topology facility passes the system CPU topology information to the scheduler, allowing it to make load balancing decisions. On machines where I/O interrupts are unevenly distributed, CPUs that are grouped together and get more I/O interrupts than others will tend to have a higher average load, creating performance issues in some cases.

Previously, CPU topology support was enabled by default. With this update, CPU topology support is disabled by default, and the kernel parameter "topology=on" has been added to allow this feature to be enabled. (BZ#475797)⁴⁶

- New kernel options can now be added using the IPL command without modifying the content of the CMS parmfile, allowing for temporary overwriting of kernel options that are already provided by the parmfile. The entire boot command line can be replaced with the VM parameter string, bypassing any kernel options from the parmfile. Furthermore, customers can create new Linux Named Saved Systems (NSS) on the CP/CMS command line. (BZ#475530)⁴⁷
- The qeth driver has been updated with HiperSockets Layer3 support for IPv6. (BZ#475572)⁴⁸ For further details on this feature, refer to the "qeth device driver for OSA-Express (QDIO) and HiperSockets" chapter in IBM's "Device Drivers, Features, and Commands" book located at: http://www.ibm.com/developerworks/linux/linux390/october2005_documentation.html
- Starting with z9 HiperSocket firmware returns the version string in a different format. This change resulted in missing mcl_level information in the qeth status message issued during online setting

of the device. The updated geth driver now correctly reads the new version string format of HiperSockets, allowing for a standardization of output format. (BZ#479881)⁴⁹

- In Red Hat Enterprise Linux 5.4, the **s390utils** package has been rebased to version 1.8.1. For a full list of features that this rebase provides, please refer to the Package Updates section of the *Technical Notes*⁵⁰. (BZ#477189)⁵¹
- In the kernel, a sysfs interface has been implemented to associate actions to shutdown triggers. For more details on this feature, refer to the "Shutdown actions" chapter in IBM's "Device Drivers, Features, and Commands" book located at: http://www.ibm.com/developerworks/linux/linux390/development_documentation.html

8. Kernel Updates

8.1. General Kernel Feature Support

- Asymmetric Logical Unit Access (ALUA) support in device-mapper-multipath has been updated, adding explicit ALUA support for Clariion storage. Earlier versions of Red Hat Enterprise Linux 5 added support for implicit ALUA (i.e. the operating system is not aware of which storage device paths have optimized performance and which have non-optimized performance). If the operating system consistently sends I/O on a non-optimized path, then the storage device may transparently make that path optimized, improving performance and causing idle paths to become non-optimized.

Red Hat Enterprise Linux 5.4 introduces *explicit* ALUA support for Clariion storage (i.e. the operating system exchanges information with the storage device and is able to select the paths that have optimized performance). (BZ#482737)⁵²

- Previously, support for raw devices in the upstream kernel was deprecated. However, this support has been returned to the kernel. Consequently, in Red Hat Enterprise Linux 5.4, support for raw devices has also been returned. Additionally, the initscripts packages have been updated, adding the previously dropped functionality of raw devices. (BZ#472891)⁵³
- KVM guest-smp tlb flushing without *mmu-notifiers* could corrupt memory as a KVM may add pages to the kernel freelist while another **vcpu** may still be writing to them through guest mode. This update adds **mmu-notifier** support to the kernel and also corrects a bug found in an earlier patch wherein *mm_struct* was grown by existing drivers and caused a failed kABI check. This bug has been corrected by using an index that resides in an unused padding hole to avoid expanding the structure size. (Bugzilla #485718)⁵⁴
- Pointer and signed arithmetic overflow wrapping has not previously been defined in the Linux kernel. This could cause **GCC** (GNU C Compiler) to assume that wrapping does not occur and attempt to optimize the arithmetic that the kernel may require for overflow testing. This update adds the *-fwrapv* variable to **GCC CFLAGS** in order to define wrapping behavior. (Bugzilla #491266)⁵⁵
- An issue of contention between processes vying for the same memory space in high end systems was recently identified by TPC-C (Transaction Processing Council) benchmarking. This update includes **fast-gup** patches which use direct IO and provide a significant (up to 9-10%) performance improvement. This update has been tested thoroughly and is used in the 5.4 kernel to improve scalability. For further information, see this *article*⁵⁶. (Bugzilla #474913)⁵⁷
- A new tunable parameter has been added to this kernel, allowing system administrators to change the maximum 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 1024 or 4MB so that a maximum of 1024 pages get written out by each iteration of `kupdate`. ([Bugzilla #479079](#)⁵⁸).

- A new option (`CONFIG_TASK_IO_ACCOUNTING=y`) has been added to kernel to assist in monitoring IO statistics per process. This assists with troubleshooting in a production environment. ([Bugzilla #461636](#)⁵⁹)
- In previous kernels, back-up processes were deteriorating DB2 server responsiveness. This was caused by `/proc/sys/vm/dirty_ratio` preventing processes writing to pagecache memory when more than half of the unmapped pagecache memory was dirty (even if `dirty_ratio` was set to 100%). A change made in this kernel update overrides this limiting behavior. Now, when the `dirty_ratio` is set to 100%, the system will no longer limit writing to pagecache memory. ([Bugzilla #295291](#)⁶⁰)
- The `rd_blocksize` option found in the previous kernel's ramdisk driver was causing data corruption when using large ramdisks under a reasonable system load. This update removes the unnecessary option and resolves the data corruption issues. ([Bugzilla #480663](#)⁶¹)
- The function `getrusage` is used to examine the resource usage of a process. It is useful in diagnosing problems and gathering data on resource usage. In instances where a process interrogated by `getrusage` was spawning child process threads, however, the results would be incorrect as `getrusage` would examine only the parent process and not interrogate its children. This update implements `rusadge_thread` to allow for accurate resource usage results in these instances. ([Bugzilla #451063](#)⁶²)
- The header `/usr/include/linux/futex.h` would previously interfere with compiling C source code files, resulting in an error. This update includes a patch which corrects problematic kernel only definitions and resolves the compiling error. ([Bugzilla #475790](#)⁶³)
- In previous kernels the kernel version was not identified in panic or oops output messages. This update adds the kernel version details to oops and panic output. ([Bugzilla #484403](#)⁶⁴)
- During release 2.6.18, the kernel was configured to provide kernel-headers for the package `glibc`. That process caused various files to be improperly marked for inclusion. The `serial_reg.h` file was incorrectly marked and not included in the `kernel_headers` rpm. This, in turn, caused problems with building other rpms. This update adds the `serial_reg.h` file and corrects the problem. ([Bugzilla #463538](#)⁶⁵)
- In some circumstances `upcrund`, the process manager in the **HP Unified Parallel C** (UPC) product, returned an `ESRCH` result and failed when calling `setpgid()` for a child process forked by a sub-thread. This update includes a patch to fix for this problem. ([Bugzilla #472433](#)⁶⁶)
- Functionality has been added to `sysrq-t` to display backtrace information about running processes. This will assist in debugging hung systems. ([Bugzilla #456588](#)⁶⁷)

8.1.1. Debugging

In Red Hat Enterprise Linux 5.4, debugging the kernel has been improved with added features for generating core dumps. Core dumps (memory snapshots) are useful for debugging system and kernel crashes. With this update, it is now possible to perform a core dump on systems that utilize hugepages. ([BZ#470411](#))⁶⁸ Additionally, kernel panic messages can now be extracted from a core dump file (vmcore) using the `makedumpfile` command. ([BZ#485308](#))⁶⁹

⁶⁸ https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=470411

8.1.2. Security

- This update increases the maximum length of the kernel key field from the arbitrary 32 character length set in previous kernels to 255 characters. ([Bugzilla #475145](#)⁷⁰)
- This kernel update addresses the security concern arising from non-root users being able to create device nodes on filesystems exported using NFS (Network File System daemon). This update returns the **CAP_MKNOD** and **CAP_LINUX_IMMUTABLE** capabilities to the user with the FSUID of 0 in the file system mask. ([Bugzilla #497272](#)⁷¹ and [Bugzilla #499076](#)⁷²)
- In keeping with Federal Information Processing Standardization 140 (FIPS140) certification requirements, this update includes:
 - Self-testing for; **ansi_cprng** ([Bugzilla #497891](#)⁷³), **ctr(aes)** mode ([Bugzilla #497888](#)⁷⁴), Hmac-sha512 ([Bugzilla #499463](#)⁷⁵), **rfc4309(ccm(aes))**. ([Bugzilla #472386](#)⁷⁶), .
 - Code to produce a signature file that GRUB performs a checksum against during the boot process. ([Bugzilla #444632](#)⁷⁷)
 - Code to change the DSA key from 512 bit to 1024 bit for module signing. ([Bugzilla #413241](#)⁷⁸)

8.2. General Platform Support

Throttling State (T-State) notification support has been added to the Advanced Configuration and Power Interface (ACPI) implementation in the kernel. Adding T-State notification enhances the use of Intel® Intelligent Power Node Manager technology for power management in data centers. ([BZ#487567](#)).⁷⁹

8.3. Driver Updates

8.3.1. Open Fabrics Enterprise Distribution (OFED) Drivers

The OpenFabrics Alliance Enterprise Distribution (OFED) is a collection of Infiniband and iWARP hardware diagnostic utilities, the Infiniband fabric management daemon, Infiniband/iWARP kernel module loader, and libraries and development packages for writing applications that use Remote Direct Memory Access (RDMA) technology. Red Hat Enterprise Linux uses the OFED software stack as its complete stack for Infiniband/iWARP/RDMA hardware support.

In Red Hat Enterprise Linux 5.4, the following portions of OFED have been updated to the upstream version 1.4.1-rc3

- Remote Direct Memory Access (RDMA) headers ([BZ#476301](#))⁸⁰
- Reliable Datagram Sockets (RDS) protocol ([BZ#477065](#)⁸¹, [BZ#506907](#))⁸²
- Sockets Direct Protocol (SDP) ([BZ#476301](#))⁸³
- SCSI RDMA Protocol (SRP) ([BZ#476301](#))⁸⁴
- IP over InfiniBand (IPoIB) ([BZ#434779](#)⁸⁵, [BZ#466086](#)⁸⁶, [BZ#506907](#))⁸⁷

⁶⁹ https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=485308

⁷⁹ https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=487567

Additionally, the following OFED drivers have been updated to the upstream version 1.4.1-rc3:

- The `cxgb3` and `iw_cxgb3` drivers for the Chelsio T3 Family of network devices ([BZ#476301](#),⁸⁸ [BZ#504906](#))⁸⁹
- The driver for mthca-based InfiniBand HCA (Host Channel Adapter) ([BZ#476301](#),⁹⁰ [BZ#506097](#))⁹¹
- `qlgc_vnic` driver ([BZ#476301](#))⁹²



Note

Red Hat closely tracks the upstream OFED code base in order to provide a maximal level of enablement for this still evolving technology. As a consequence, Red Hat can only preserve API/ABI compatibility across minor releases to the degree that the upstream project does. This is an exception from the general practice in the development of Red Hat Enterprise Linux.

8.3.2. General Driver Updates

- The **iw5400** driver for Intel 5400 class Memory Controllers has been updated with added support for Error Detection And Correction (EDAC). ([BZ#462895](#))⁹³
- the **i2c** driver for the iic-bus interface has been updated, adding support for the AMD SB800 Family of products.
- The **i2c-piix4** driver has been updated with support for the Broadcom HT1100 chipset. ([BZ#474240](#))⁹⁴
- The `hpilo` driver has been updated. ([BZ#488964](#))⁹⁵.
- The **dm9601** for Davicom Ethernet Adaptors has been updated.

8.3.3. Network Driver Updates

- Red Hat is no longer able to responsibly support the **ipw3945** wireless laptop network driver to the full levels of service, as the upstream company responsible for that driver has suspended maintenance efforts on the driver.

Red Hat, following the guidance of the company responsible for the driver has migrated over to the **iw13945** driver (first introduced in Red Hat Enterprise Linux 5.3). Ongoing maintenance of driver issues associated with the 3945 hardware will be provided on the **iw13945** driver. All fresh/full installations of Red Hat Enterprise Linux 5.4 on systems with 3945 hardware will install the `iw13945` driver.

Red Hat continues to provide the **ipw3945** driver in Red Hat Enterprise Linux 5.4 as a migration convenience. It is recommended that customers transition to the **iw13945** driver. Refer to the [Red Hat Knowledgebase](#)⁹⁶ for details on how to transition if you are doing an update install from a system that had its full installation prior to Red Hat Enterprise Linux 5.4.

Customers who report issues with the **ipw3945** driver may be referred to migrate to the **iw13945** driver for resolution. For example, there are known issues connecting to VPN concentrators present in the **ipw3945** driver which are resolved in the **iw13945** driver.



Note

the use of the `iwl3945` driver requires installing the `iwl3945-firmware` package from the supplementary RHN channel. Package version 15.28.2.8-2 or later is required. (ie `iwl3945-firmware-15.28.2.8-2` or later)

Additionally, the device name, for example as it appears in the output of the ``ifconfig`` command appears as `eth0` (or `eth1`) when using the `ipw3945` driver, when using the `iwl3945` driver it is identified as `wlan0`.

- the bonding driver has been updated to the latest upstream version. This update, however has introduced symbol/ipv6 module dependency capabilities. Therefore, if IPv6 has been previously disabled (by inserting the `install ipv6 /bin/false` line in the `/etc/modprobe.conf` file) an upgrade to the bonding driver in 5.4 will result in the bonding kernel module failing to load. The `install ipv6 /bin/false` line needs to be replaced with `options ipv6 "disable=1` for the module to load properly.
- The drivers in the kernel for Intel® I/O Acceleration Technology (Intel® I/OAT) have been updated to version 2.6.24.[\(BZ#436048\)](#)⁹⁷.
- the `igb` driver for Intel® Gigabit Ethernet Adapters has been updated to version 1.3.16-k2. This update also enables GRO support for the `igb` driver. [\(BZ#484102,](#) ⁹⁸ [BZ#474881,](#) ⁹⁹ [BZ#499347\)](#).¹⁰⁰
- The `igbvf` driver has been updated, providing Virtual Function support for Intel 82576 Gigabit Ethernet Controllers. [\(BZ#480524\)](#)¹⁰¹
- the `ixgbe` driver for Intel 10 Gigabit PBetaCI Express network devices has been updated to version 2.0.8-k2. This update also enables GRO support for the `ixgbe` driver. [\(BZ#472547,](#) ¹⁰² [BZ#499347\)](#).¹⁰³
- The `bnx2` driver for Broadcom NetXtreme II network devices has been updated to version 1.9.3 [\(BZ#475567](#)¹⁰⁴ [BZ#476897](#)¹⁰⁵ [BZ#489519\)](#)¹⁰⁶
- The `tg3` driver for Broadcom Tigon3 ethernet devices has been updated to version 3.96. [\(BZ#481715,](#) ¹⁰⁷ [BZ#469772\)](#).¹⁰⁸ This driver update adds support for 5785F and 50610M devices. [\(BZ#506205\)](#)¹⁰⁹
- The `cnic` driver has been added, providing Internet Small Computer System Interface (iSCSI) support for `bnx2` network devices. [\(BZ#441979\)](#)¹¹⁰.
- The `bnx2x` driver for Broadcom Everest network devices has been updated to version 1.48.105.[\(BZ#475481\)](#)¹¹¹.
- The `bnx2i` driver has been added, providing iSCSI support for `bnx2x` network devices. [\(BZ#441979\)](#)¹¹².
- The `cxgb3` driver for the Chelsio T3 Family of network devices has been updated, enabling iSCSI TCP Offload Engines (TOE) and Generic Receive Offload (GRO) support. [\(BZ#439518](#)¹¹³, [BZ#499347\)](#)¹¹⁴
- The `forcedeth` ethernet driver for NVIDIA nForce devices has been updated to version 0.62. [\(BZ#479740\)](#).¹¹⁵

- The sky2 driver for ethernet controllers using the Marvell Yukon 2 chipset has been updated. ([BZ#484712](#)).¹¹⁶
- The enic driver for Cisco 10G ethernet devices has been updated to version 1.0.0.933. ([BZ#484824](#))¹¹⁷
- The e1000e driver for Intel PRO/1000 ethernet devices has been updated to the upstream version 1.0.2-k2. ([BZ#480241](#))¹¹⁸
- The be2net driver for Emulex Tiger Shark converged network adapters has been added as a Technology Preview.

8.3.4. Storage Driver Updates

- The **bnx2** driver now supports iSCSI. The **bnx2i** driver will access the **bnx2** driver through the **cnic** module to provide iSCSI offload support. To manage **bnx2i**, use the **iscsi-initiator-utils** package. For instructions on **bnx2i** configuration, please refer to *section 5.1.2* of the `/usr/share/docs/iscsi-initiator-utils-<version>/README` file. ([BZ#441979](#)¹¹⁹ and [BZ#441979](#)¹²⁰)

Note that the **bnx2i** version included in this release does not support IPv6.

- The **md** driver has been updated to provide support for *bitmap merging*. This feature eliminates the need for full resync when performing data replication. ([BZ#481226](#))¹²¹
- The **scsi** layer in this release features the following updates:
 - The **scsi** driver now includes the upstream **scsi_dh_alua** module. This adds explicit *asymmetric logical unit access* (ALUA) support with this release. To utilize the **scsi_dh_alua** module when using **dm-multipath**, specify **alua** as the **hardware_handler** type in **multipath.conf**. ([BZ#482737](#))¹²²

Note that for *EMC Clariion* devices, using only **scsi_dh_alua** or **dm-emc** alone is supported. Using both **scsi_dh_alua** and **dm-emc** is not supported.

- The **rdac_dev_list** structure now includes **md3000** and **md3000i** entries. This allows users to benefit from the advantages provided by the **iscsi_dh_rdac** module. ([BZ#487293](#))¹²³
- A bug that caused iSCSI iBFT installations to panic during disk formatting is now fixed. ([BZ#436791](#))¹²⁴
- A bug in the **iscsi_r2t_rsp_struct** that caused kernel panics during iSCSI failovers in some multipathed environments is now fixed. ([BZ#484455](#))¹²⁵
- The **cxgb3** driver has been updated to apply several upstream fixes and provide support for iSCSI TOE devices. ([BZ#439518](#))¹²⁶

Note that the **cxgb3i** version included in this release does not support IPv6.

- This release includes the new **mpt2sas** driver. This driver supports the SAS-2 family of adapters from LSI Logic. SAS-2 increases the maximum data transfer rate from 3Gb/s to 6Gb/s.

The **mpt2sas** driver is located in the **drivers/scsi/mpt2sas** directory, as opposed to the older **mpt** drivers that are located in **drivers/message/fusion** directory. ([BZ#475665](#))¹²⁷

- The **aacraid** driver has now been updated to version 1.1.5-2461. This update applies several upstream fixes for bugs affecting queued scans, controller boot problems, and other issues. ([BZ#475559](#)¹²⁸)
- The **aic7xxx** driver now features an increased maximum I/O size. This allows supported devices (such as SCSI tape devices) to perform writes with larger buffers.
- The **cciss** driver has been updated to apply upstream fixes for bugs affecting memory BAR discovery, the **rebuild_lun_table** and the MSA2012 scan thread. This update also applies several configuration changes to **cciss**.
- The **fnic** driver has been updated to version 1.0.0.1039. This applies several upstream bug fixes, updates the **libfc** and **fcoe** modules, and adds a new module parameter that controls debug logging at runtime. ([BZ#484438](#)¹²⁹)
- The **ipr** driver now supports MSI-X interrupts. ([BZ#475717](#)¹³⁰)
- The **lpfc** driver has been updated to version 8.2.0.48. This enables hardware support for upcoming OEM programs. In addition, this update also applies the following bug fixes (among others): ([BZ#476738](#)¹³¹ and [BZ#509010](#)¹³²)

- Virtualized fibre-channel switches are now supported.
- Polling for error attention interrupts are now available.
- A bug that caused memory leaks in **vport create** and **delete loop** is now resolved.

With this update, the **lpfc** driver now also supports *HBAAnyware 4.1* and *OneConnect UCNA*. ([BZ#498524](#)¹³³)

- The **MPT fusion** driver is now updated to version 3.04.07rh v2. This applies several bug fixes, including: ([BZ#475455](#)¹³⁴)
 - An **MPT fusion** driver bug that prevented the system from booting with the PAE kernel is now fixed.
 - Controllers are now set to **READY_STATE** when the driver unloads.
 - The **mptsas** driver now issues **TUR** (Test Unit Ready) and **Report LUN** commands before adding a device to the transport layer.

In addition, a patch that unexpectedly caused **mptctl_ioctl()** to issue numerous yet benign kernel error messages is now reverted. With this release, **mptctl_ioctl()** no longer issues these kernel error messages.

- The **megaraid_sas** driver is now updated to version 4.08-RH1. This update applies the following upstream enhancements and fixes (among others):([BZ#475574](#)¹³⁵)
 - This update adds a polling mode to the driver.
 - A bug affecting supported tape drives is now fixed. With this release, the **pthru** timeout value is now set to the O/S layer timeout value for commands sent to tape drives.

- The **mvsas** driver is now updated to version 0.5.4. This applies several fixes and enhancements from upstream, and adds support for *Marvell RAID* bus controllers MV64460, MV64461, and MV64462. ([BZ#485126](#)¹³⁶)
- The **qla2xxx** driver has been updated to version 8.03.00.10.05.04-k, and now supports *Fibre Channel over Convergence Enhanced Ethernet* adapters. With this release, **qla2xxx** also applies several bug fixes from upstream, including: ([BZ#471900](#)¹³⁷, [BZ#480204](#)¹³⁸, [BZ#495092](#)¹³⁹, and [BZ#495094](#)¹⁴⁰)
 - Discrepancies detected during **OVERRUN** handling on 4GB and 8GB adapters are now corrected.
 - All **vports** are now alerted of any asynchronous events.
 - A bug that caused kernel panics with the *QLogic 2472* card is now fixed.
 - The **stop_firmware** command is no longer retried if the first attempt results in a times out.
 - The sector mask value is no longer based on the fixed **optrom** size.
 - A bug that caused frequent path failures during I/O on multipathed devices is now fixed. ([BZ#244967](#)¹⁴¹)
 - The driver source code is now kABI-compliant.
 - **dcbx** pointers are now set to **NULL** after freeing memory.

In addition to these updates, the **qla24xx** and **qla25xx** firmwares included in the **qla2xxx** driver are now updated to version 4.04.09.

- The **qla4xxx** driver now features improved driver fault recovery. This update fixes a bug in the driver that prevented adapter recovery if there were outstanding commands detected on the host adapter. ([BZ#497478](#)¹⁴²)
- This release includes the new **qlge** driver. This driver adds ethernet support for *QLogic FCoE* 10GB adapters. ([BZ#479288](#)¹⁴³)

9. Technology Previews

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

The following Technology Previews are new or enhanced in Red Hat Enterprise Linux 5.4. For detailed information on the Technology Previews in Red Hat Enterprise Linux 5.4, refer to the Technology Previews section of the 5.4 Technical Notes located at <http://www.redhat.com/docs/manuals/enterprise/>

- New Emulex Tiger Shark converged network adapter driver. Refer to [Section 8.3.3, “Network Driver Updates”](#) for more information.
- New Systemtap Tracepoints. Refer to [Section 8.1.1, “Debugging”](#) for more information.
- New Gnu Compiler Collection version 4.4. Refer to [Section 6, “Tools Updates”](#) for more information.

- Updated ext4 support. Refer to [Section 4, “Filesystems and Storage updates”](#) for more information.
- glibc new MALLOC behaviour. Refer to [Section 6, “Tools Updates”](#) for more information.
- samba3x and ctdb on the x86_64 platform. Refer to [Section 4, “Filesystems and Storage updates”](#) for more information.
- Clustering with KVM hypervisor. Refer to [Section 2, “Clustering Updates”](#) for more information.
- The **fence_virsh** fence agent is provided in this release of Red Hat Enterprise Linux as a Technology Preview. Refer to [Section 2.1, “Fencing Improvements”](#) for more information.
- Stream Control Transmission Protocol (SCTP) support for the Cluster Suite and the redundant ring configuration of clusters are considered Technology Previews in Red Hat Enterprise Linux 5.4. ([BZ#457956](#))¹⁴⁴
- USB passthrough with the KVM hypervisor is considered to be a Technology Preview for the 5.4 release.
- In Red Hat Enterprise Linux 5.4, Logical Partition (LPAR) support for the cluster-suite is considered a Technology Preview. ([BZ#480662](#))¹⁴⁵, ([BZ#498083](#))¹⁴⁶
- The **xorg-x11-qxl-drv** X11 video driver for the qemu QXL video accelerator has been added to 5.4 as a Technology Preview. ([BZ#489806](#))¹⁴⁷

A. Revision History

Revision 1.0 Wed Sep 02 2009

Ryan Lerch rlerch@redhat.com

Initial version of the online version of the Red Hat Enterprise Linux 5.4 Release Notes

