

Red Hat Enterprise Linux 6

6.2 Release Notes

Release Notes for Red Hat Enterprise Linux 6.2



Red Hat Enterprise Linux 6 6.2 Release Notes

Release Notes for Red Hat Enterprise Linux 6.2

Edition 2

Copyright © 2011 Red Hat, Inc.

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

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

Red Hat Enterprise Linux minor releases are an aggregation of individual enhancement, security and bug fix errata. The Red Hat Enterprise Linux 6.2 Release Notes documents the major changes made to the Red Hat Enterprise Linux 6 operating system and its accompanying applications for this minor release. Detailed notes on all changes in this minor release are available in the [Technical Notes](#)¹.

¹ https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/6.2_Technical_Notes/index.html

Preface	v
1. Hardware Support	1
2. Installation	3
3. Kernel	5
4. Resource Management	13
5. Device Drivers	15
6. Storage	19
7. File System	21
8. Networking	23
9. Authentication and Interoperability	25
10. Entitlement	27
11. Security, Standards and Certification	29
12. Compiler and Tools	31
13. Clustering	33
14. High Availability	35
15. Virtualization	37
16. Graphics	41
17. General Updates	43
A. Component Versions	45
B. Revision History	47

Preface

The Release Notes provide high level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 6.2. For detailed documentation on all changes to Red Hat Enterprise Linux for the 6.2 update, refer to the [Technical Notes](#)¹.



Note

Refer to the [Online Release Notes](#)² for the most up-to-date version of the Red Hat Enterprise Linux 6.2 Release Notes.

¹ https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/6.2_Technical_Notes/index.html

² https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/6.2_Release_Notes/index.html

Hardware Support

biosdevname

The *biosdevname* package has been upgraded to version 0.3.8, providing the **--smbios** and **--nopirq** command line parameters. With these command line parameters, source code patches, which removed these codepaths, can be removed from the build process.

Installation

Compression format for the `initrd.img` file

In Red Hat Enterprise Linux 6.2, the compression format used for the `initrd.img` file has changed. The image is now compressed using LZMA rather than Gzip.

To decompress an image, use the `xz -d` command. For example:

```
~]# xz -dc initrd.img | cpio -id
```

To compress an image, use the `xz -9 --format=lzma` command. For example:

```
~]# find . | cpio -c -o | xz -9 --format=lzma > initrd.img
```

Support for device identification using WWIDs during installation

Fibre Channel and Serial Attach SCSI (SAS) devices can be now specified by a World Wide Name (WWN) or a World Wide Identifier (WWID) for unattended installations. WWN is part of the IEEE standard which makes it easier to identify storage devices during installation for users utilizing Storage Area Networks (SAN) and other advanced network topologies. When a storage device is attached to a server using multiple physical paths for redundancy or improved performance, WWN for any of these paths is sufficient to identify the device.

Initial ramdisk file

The initial ramdisk file on 64-bit PowerPC and 64-bit IBM POWER Series systems is now named `initrd.img`. In previous releases, it was named `ramdisk.image.gz`.

Static IPv6 address support for network installation

In Red Hat Enterprise Linux 6.2, a static IPv6 address can be specified for the `ipv6` boot option for network installations. The specified address must be of the following form:

```
<IPv6 address>[/<prefix length>]
```

An example of a valid IPv6 address would then be `3ffe:ffff:0:1::1/128`. If the prefix is omitted, the value of 64 is assumed. Specifying a static IPv6 address for the `ipv6` boot option complements the already existing `dhcp` and `auto` parameters that can be specified for the `ipv6` boot option.

Kernel

The kernel shipped in Red Hat Enterprise Linux 6.2 includes several hundred bug fixes for and enhancements to the Linux kernel. For details concerning every bug fixed in and every enhancement added to the kernel for this release, refer to the kernel section of the [Red Hat Enterprise Linux 6.2 Technical Notes](#).¹

Using open-iscsi to manage the qla4xxx discovery and login process

Prior to Red Hat Enterprise Linux 6.2, the qla4xxx adapter firmware managed discovery and login to iSCSI targets. A new feature in Red Hat Enterprise Linux 6.2 allows you to use **open-iscsi** to manage the qla4xxx discovery and login process. This can result in a more uniform management process.

This new feature is enabled by default. The qla4xxx iSCSI firmware settings are accessible via:

```
~]# iscsiadm -m fw
```

This feature may be disabled by setting the module `ql4xdisablesysfsboot=1` parameter as follows:

1. Set the parameter in the `/etc/modprobe.d` file:

```
~]# echo "options qla4xxx ql4xdisablesysfsboot=1" >> /etc/modprobe.d/qla4xxx.conf
```

2. Reload the qla4xxx module either by executing the following set of commands:

```
~]# rmmod qla4xxx
~]# modprobe qla4xxx
```

or, if you are booted off the qla4xxx device, by rebooting your system.

When booted off a qla4xxx device, upgrading from Red Hat Enterprise Linux 6.1 to Red Hat Enterprise Linux 6.2 will cause the system to fail to boot up with the new kernel. For more information on this known issue, refer to the [Technical Notes](#).²

kexec kdump support on additional file systems

Kdump (a kexec-based crash dumping mechanism) now supports dumping of the core on the following file systems on Red Hat Enterprise Linux 6:

- Btrfs (Note that this file system is a Technology Preview)
- ext4
- XFS (Note that XFS is a layer product and must be installed to enable this feature)

¹ https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/6.2_Technical_Notes/index.html

² https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/6.2_Technical_Notes/kernel_issues.html

pkgtemp merged with coretemp

The `pkgtemp` module has been merged with the `coretemp` module. The `pkgtemp` module is now deprecated. The `coretemp` module now supports all the features it previously did plus the features that were supported by the `pkgtemp` module.

The `coretemp` previously only provided per core temperatures, while the `pkgtemp` module provided the temperatures of the CPU package. In Red Hat Enterprise Linux 6.2, the `coretemp` module allows you to read the temperatures for of the cores, the uncore, and the package.

It is advisable to adjust any scripts using either of these modules.

Lockless dispatching of SCSI driver queuecommand functions

In Red Hat Enterprise Linux 6.2, the SCSI midlayer supports optional lockless dispatching of SCSI driver queuecommand functions.

This is a backport of the upstream SCSI lock pushdown commit. The backport retains binary compatibility with Red Hat Enterprise Linux 6.0 and Red Hat Enterprise Linux 6.1. Retaining binary compatibility requires divergence from the equivalent upstream SCSI lock pushdown mechanism.

A previously unused flag in the `scsi_host_template` structure is used by SCSI drivers to indicate to the SCSI midlayer that driver queuecommand will be dispatched without the SCSI host bus lock held.

The default behavior is that the `Scsi_Host` lock will be held during a driver queuecommand dispatch. Setting the `scsi_host_template` lockless bit prior to `scsi_host_alloc` will cause the driver queuecommand function to be dispatched without the `Scsi_Host` lock being held. In such a case, the responsibility for any lock protection required is *pushed down* into the driver queuecommand code path.

SCSI Drivers updated to use lockless queuecommand in Red Hat Enterprise Linux 6.2 are listed below:

- `iscsi_iser`
- `be2iscsi`
- `bnx2fc`
- `bnx2i`
- `cxgb3i`
- `cxgb4i`
- `fcoe` (software fcoe)
- `qla2xxx`
- `qla4xxx`

Support for Fiber Channel over Ethernet (FCoE) target mode

Red Hat Enterprise Linux 6.2 includes support for Fiber Channel over Ethernet (FCoE) target mode, as a *Technology Preview*. This kernel feature is configurable via `targetadmin`, supplied by the `fcoe-target-utils` package. FCoE is designed to be used on a network supporting Data Center Bridging (DCB). Further details are available in the `dcbtool(8)` and `targetadmin(8)` man pages.



Important

This feature uses the new SCSI target layer, which falls under this Technology Preview, and should not be used independently from the FCoE target support. This package contains the AGPL license.

Support for the `crashkernel=auto` boot parameter

In Red Hat Enterprise Linux 6.1, with BZ#[605786](https://bugzilla.redhat.com/show_bug.cgi?id=605786)³, the `crashkernel=auto` boot parameter was deprecated. However, in Red Hat Enterprise Linux 6.2, support for `crashkernel=auto` is continued on all Red Hat Enterprise Linux 6 systems.

Support for MD RAID in user space

The `mdadm` and `mdmon` utilities have been updated to support Array Auto-Rebuild, RAID Level Migrations, RAID 5 support limitation, and SAS-SATA drive roaming.

Flush request merge

Red Hat Enterprise Linux 6.2 supports merging of flush requests to assist devices which are slow to perform a flush.

UV2 Hub Support

Red Hat Enterprise Linux 6.2 adds UV2 Hub support. UV2 is the UVhub chip that is the successor to the current UV1 hub chip. UV2 uses the HARP hub chip that is currently in development. UV2 provides support for new Intel sockets. It provides new features to improve performance. UV2 is being designed to support 64 TB of memory in an SSI. Additionally, the node controller MMRs have been updated for UV systems.

`acpi_rsdp` boot parameter

Red Hat Enterprise Linux 6.2 introduces the `acpi_rsdp` boot parameter for `kdump` to pass an ACPI RSDP address, so that the `kdump` kernel can boot without EFI (Extensible Firmware Interface).

QETH driver improvements

The following enhancements have been added to the QETH network device driver:

- Support for `af_iucv` HiperSockets transport
- Support for forced signal adapter indications
- Support for asynchronous delivery of storage blocks
- New Ethernet Protocol ID added to the `if_ether` module

³ https://bugzilla.redhat.com/show_bug.cgi?id=605786

CPACF algorithms

Support for the new CPACF (CP Assist for Cryptographic Function) algorithms, supported by IBM zEnterprise 196, has been added. The new hardware accelerated algorithms are:

- CTR mode for AES
- CTR mode for DES and 3DES
- XTS mode for AES with key lengths of 128 and 256 bits
- GHASH message digest for GCM mode

Red Hat Enterprise Linux 6.2 supports conditional resource-reallocation through the *pci=realloc* kernel parameter. This feature provides an interim solution for adding a dynamic reallocate pci resource without causing any regressions. It disables dynamic reallocation by default, but adds the ability to enable it through the *pci=realloc* kernel command line parameter.

PCI improvements

Dynamic reallocation is disabled by default. It can be enabled with the **pci=realloc** kernel command line parameter. In addition, bridge resources have been updated to provide larger ranges in the **PCI assign unassigned** call.

SMEP

Red Hat Enterprise Linux 6.2 enables SMEP (Supervision Mode Execution Protection) in the kernel. SMEP provides an enforcement mechanism, allowing the system to set a requirement that is not intended to be executed from user pages while in the supervisor mode. This requirement is then enforced by the CPU. This feature is able to prevent all attacks irrespective of the vulnerability in the system code that are executed from user mode pages while the CPU is in the supervisor mode.

Enhanced fast string instructions

Support for enhanced fast string **REP MOVSB/STOESB** instructions for the latest Intel platform has been added.

USB 3.0 xHCI

The USB 3.0 xHCI host side driver has been updated to add split-hub support, allowing the xHCI host controller to act as an external USB 3.0 hub by registering a USB 3.0 roothub and a USB 2.0 roothub.

ACPI, APEI, and EINJ parameter support

The ACPI, APEI, and EINJ parameter support is now disabled by default.

pstore

Red Hat Enterprise Linux 6.2 adds support for **pstore**—a file system interface for platform dependent persistent storage.

PCIe AER error information printing

Support for printk based APEI (ACPI Platform Error Interface) hardware error reporting has been added, providing a way to unify errors from various sources and send them to the system console.

ioatdma driver

The ioatdma driver (dma engine driver) has been updated to support Intel processors with a dma engine.

8250 PCI serial driver

Support for the Digi/IBM PCIe 2-port Async EIA-232 Adapter has been added to the 8250 PCI serial driver. Additionally, EEH (Enhanced Error Handling) support for the Digi/IBM PCIe 2-port Async EIA-232 Adapter has been added to the 8250 PCI serial driver.

ARI support

ARI (Alternative Routing- ID Interpretation) support, a PCIe v2 feature, has been to Red Hat Enterprise Linux 6.2.

PCIe OBFF

PCIe OBFF (Optimized Buffer Flush/Fill) enable/disable support has been added for Intel's latest platform. OBFF provides devices with information on interrupts and memory activity and their potentially reduced power impact, ultimately improving energy efficiency.

Capture oops/panic reports to NVRAM

In Red Hat Enterprise Linux 6.2, the kernel is enabled to capture kernel oops/panic reports from the dmesg buffer into NVRAM on PowerPC architectures.

MXM driver

The MXM driver, responsible for handling graphics switching on NVIDIA platforms, has been backported to Red Hat Enterprise Linux 6.2.

Page coalescing

Red Hat Enterprise Linux 6.2 introduces page coalescing, a feature on IBM Power servers which allows for coalescing identical pages between logical partitions.

L3 cache partitioning

Support for L3 Cache Partitioning has been added to the latest AMD family CPUs.

thinkpad_acpi module

The thinkpad_acpi module has been updated to add support for new ThinkPad models.

C-State support

Latest Intel processor C-State support has been added to **intel_idle**.

IOMMU warnings

Red Hat Enterprise Linux 6.2 now displays warnings for IOMMU (Input/Output Memory Management Unit) on AMD systems.

Logging to dmesg during boot

Logging of board, system, and BIOS information to dmesg during boot has been added.

IBM PowerPC support

`cputable` entries have been added to the kernel, providing support for the latest IBM PowerPC processor family.

VPHN

The VPHN (Virtual Processor Home Node) feature has been disabled on IBM System p.

Drivers supported by latest Intel chipset

The following drivers are now supported by the latest Intel chipset:

- `i2c-i801` SMBus driver
- `ahci` AHCI-mode SATA
- `ata_piix` IDE-mode SATA driver
- TCO Watchdog driver
- LPC Controller driver

exec-shield

On IBM PowerPC systems, the `exec-shield` value in `sysctl` or in the `/proc/sys/kernel/exec-shield` parameter is no longer enforced.

kdump on PPC64

Additional checks and fixes have been added to support `kdump` on 64-bit PowerPC and 64-bit IBM POWER Series systems.

UV MMTIMER module

The UV MMTIMER module (`uv_mmtimer`) has been enabled on SGI platforms. The `uv_mmtimer` module allows direct userland access to the UV system's real time clock which is synchronized across all hubs.

IB700 module

Support for the IB700 module has been added in Red Hat Enterprise Linux 6.2

Override PCIe AER Mask Registers

The `aer_mask_override` module parameter has been added, providing a way to override the corrected or uncorrected masks for a PCI device. The mask will have the bit corresponding to the status passed into the `aer_inject()` function.

USB 3.0 host controller support on PPC64

USB 3.0 host controller support has been added to 64-bit PowerPC and 64-bit IBM POWER Series systems.

OOM killer improvements

An improved upstreamed OOM (Out of Memory) killer implementation has been backported to Red Hat Enterprise Linux 6.2. The improvements include:

-
- Processes which are about to exit are preferred by the OOM killer.
 - The OOM kill process also kills the children of the selected processes.
 - Heuristic have been added to kill the forkbomb processes.

The `oom_score_adj` **/proc** tunable parameter adds the value stored in each process's `oom_score_adj` variable, which can be adjusted via **/proc**. This allows for an adjustment of each process's attractiveness to the OOM killer in user space; setting it to `-1000` will disable OOM kills entirely, while setting to `+1000` is marking this process as OOM's primary kill target.

For more information on the new implementation, refer to <http://lwn.net/Articles/391222/>.

zram driver

Red Hat Enterprise Linux 6.2 provides an updated zram driver (creates generic RAM based compressed block devices).

taskstat utility

In Red Hat Enterprise Linux 6.2, the **taskstat** utility (prints ASET tasks status) in the kernel has been enhanced by providing microsecond CPU time granularity for the **top** utility to use.

perf utility

Red Hat Enterprise Linux 6.2 updates the **perf** utility to upstream version 3.1 along with the kernel upgrade to v 3.1. Refer to BZ#[725524](#)⁴ for newly supported kernel features provided by the **perf** utility. The updated version of the **perf** utility includes:

- Added cgroup support
- Added handling of **/proc/sys/kernel/kptr_restrict**
- Added more cache-miss percentage printouts
- Added the **-d -d** and **-d -d -d** options to show more CPU events
- Added the **--sync/-S** option
- Added support for the `PERF_TYPE_RAW` parameter
- Added more documentation about the **-f/--fields** option
- The `python-perf` package has been added for python binding support.

OProfile support

Red Hat Enterprise Linux 6.2 adds OProfile support for the latest Intel processors.

IRQ counting

The number of interrupt requests (IRQ) is now counted in a *sum of all irq* counter, reducing the cost of the look-up in the **/proc/stat** file.

⁴ https://bugzilla.redhat.com/show_bug.cgi?id=725524

Scheduling improvement

Red Hat Enterprise Linux 6.2 introduces a scheduling improvement where a hint is provided to the scheduler on the next buddy hint on sleep and preempt path. This hint/enhancement helps the workload of multiple tasks in multiple task groups.

Transparent Huge Page improvement

In Red Hat Enterprise Linux 6.2, Transparent Huge Pages are now supported in several places of the kernel:

- The system calls of **mremap**, **mincore**, and **mprotect**
- **/proc** tunable parameters: **/proc/<pid>/smaps** and **/proc/vmstat**

Additionally, Transparent Huge Pages add some compaction improvements.

XTS AES256 self-tests

Red Hat Enterprise Linux 6.2 adds XTS (XEX-based Tweaked CodeBook) AES256 self-tests to meet the FIPS-140 requirements.

SELinux netfilter packet drops

Previously, the SELinux **netfilter** hooks returned *NF_DROP* if they dropped a packet. In Red Hat Enterprise Linux 6.2, a drop in the **netfilter** hooks is signaled as a permanent fatal error and is not transient. By doing this, the error is passed back up the stack, and in some places and applications will get a faster interaction that something went wrong.

LSM hook

In Red Hat Enterprise Linux 6.2, the remount mount options (**mount -o remount**) are passed to a new LSM hook.

Default mode for UEFI systems

Red Hat Enterprise Linux 6.0 and 6.1 defaulted to running UEFI systems in a physical addressing mode. Red Hat Enterprise Linux 6.2 defaults to running UEFI systems in a virtual addressing mode. The previous behavior may be obtained by passing the *physefi* kernel parameter.

Default method for kdumping over SSH

In Red Hat Enterprise Linux 6, the default `core_collector` method for kdumping the core over SSH has been changed from *scp* to *makedumpfile*, which helps shrink the size of the core file when copying over the network link, resulting in faster copying.

If you require the old vmcore full size core file, specify the following in the `/etc/kdump.conf` file:

```
core_collector /usr/bin/scp
```

Resource Management

Cgroups CPU ceiling enforcement

The Completely Fair Scheduler (CFS) in the Linux kernel is a proportional share scheduler which divides the CPU time proportionately between groups of tasks depending on the priority/weight of the task or shares assigned to groups of tasks. In CFS, a task group can get more than its share of CPU if there are enough idle CPU cycles available in the system, due to the work conserving nature of the scheduler.

However there are enterprise scenarios listed below, where giving more than the desired CPU share to a task group is not acceptable:

Pay-per-use

In enterprise systems that cater to multiple customers, cloud service providers need to assign a fixed amount of CPU time to the virtual guest based on the service level.

Service level guarantees

Customer demands a percentage of CPU resource without service interruptions for each virtual guest.

In these scenarios, the scheduler needs to put a hard stop on the CPU resource consumption of a task group if it exceeds a preset limit. This is usually achieved by throttling the task group when it fully consumes its allocated CPU time.

The cgroups CPU ceiling enforcement is considered a very important addition to the Red Hat Enterprise Linux feature repertoire, for the use case listed above. The CPU ceiling enforcement is provided by the Credit Scheduler in Xen, and also in the VMware ESX scheduler.

Cgroups CPU controller scalability improvement on SMP systems

Red Hat Enterprise Linux 6 enabled cgroups out of the box, and **libvirt** created a cgroups per guest model. On large SMP systems, an increase in the number of cgroups, worsened the performance. However, in Red Hat Enterprise Linux 6.2, the cgroups CPU scalability has been significantly improved, making it possible to create and run several hundreds of cgroups at once with no performance implications.

In addition to the scalability improvement, a **/proc** tunable parameter, **dd sysctl_sched_shares_window**, has been added, with the default is set to 10 ms.

Cgroups I/O controller performance improvement

The cgroups I/O controller design has improved to reduce the usage of locks inside the I/O controller, resulting in improved performance. Also, the I/O controller now supports per cgroup statistics.

Cgroups memory controller performance improvement

Red Hat Enterprise Linux 6.2 introduces the memory usage overhead improvement on the memory controller by reducing allocation overhead for `page_cgroup` array by 37%. Additionally, the direct `page_cgroup-to-page` pointer has been removed, improving performance of the memory controller.

Default value for the CFQ group_isolation variable

The default for CFQ's `group_isolation` variable has been changed from 0 to 1 (`/sys/block/<device>/queue/iosched/group_isolation`). After various testing and numerous user reports,

it was found that having default 1 is more useful. When set to 0, all random I/O queues become part of the root cgroup and not the actual cgroup which the application is part of. Consequently, this leads to no service differentiation for applications.



Further Reading

For more information on resource management and control groups, refer to the [Red Hat Enterprise Linux 6.2 Resource Management Guide](https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Resource_Management_Guide/index.html)¹.

¹ https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Resource_Management_Guide/index.html

Device Drivers

Emulex lpfc driver

The default interrupt configuration for the Emulex LPFC FC/FCoE driver has changed from INT-X to MSI-X. This is reflected by the `lpfc_use_msi` module parameter (in `/sys/class/scsi_host/host#/lpfc_use_msi`) being set to 2 by default, instead of the previous 0. For more information on this change, refer to [Red Hat Enterprise Linux 6.2 Technical Notes](#)¹.

Storage drivers

- The `lpfc` driver for Emulex Fibre Channel Host Bus Adapters has been updated to version 8.3.5.45.2p.
- The `mptfusion` driver has been updated to version 3.4.19.
- The `bnx2fc` for the Broadcom Netxtreme II 57712 chip has been updated to version of 1.0.4.
- The `qla2xxx` driver for QLogic Fibre Channel HBAs has been updated to version to 8.03.07.05.06.2-k.
- The `megaraid` driver has been updated to version v5.38.
- The `arcmsr` driver for Areca RAID controllers has been updated.
- The `beiscsi` driver has been updated to version 2.103.298.0.
- The `ipr` driver for IBM Power Linux RAID SCSI HBAs has been updated to version 2.5.2.
- The `cciss` driver has been updated to provide a fix for `cciss` driver **kdump** failures.
- The `hpsa` driver has been updated to provide a fix for `hpsa` driver **kdump** failures.
- The `bnx2i` driver for Broadcom NetXtreme II iSCSI has been updated to version 2.7.0.3 to support the 578xx family of Multi-Port Single-Chip 10G Ethernet Converged Controllers.
- The `mpt2sas` driver has been updated to version 09.101.00.00.
- The Brocade BFA FC SCSI driver (`bfa` driver) has been updated to version 2.3.2.4.
- The `be2iscsi` driver for ServerEngines BladeEngine 2 Open iSCSI devices has been updated to version 4.0.160r.
- The `ata_generic` driver has been updated to add Intel IDE-R ATA support.
- The `isci` driver has been updated to version 2.6.40-rc.
- The `libfc`, `libfcoc`, and `fcoc` drivers have been updated.
- The `qib` driver TrueScale HCAs has been updated.
- The `libata` module has been updated to include improved error handling.

¹ https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/6.2_Technical_Notes/kernel_issues.html

- The md driver has been updated to include **dm-raid** target, which provides improved RAID capabilities through a DM interface. The **dm-raid** code is currently marked as a Technology Preview.
- Device Mapper support has been updated to upstream version 3.1+.
- Application support for the qla4xxx using bsg/netlink interfaces has been added.
- The DIF/DIX kernel code has been updated to the latest upstream version, affecting scsi, block, and dm/md.

Network drivers

- The netxen driver for NetXen Multi port (1/10) Gigabit Network devices to version 4.0.75.
- The vmxnet3 driver has been updated.
- The bnx2x driver has been updated to version v1.70.
- The be2net driver for ServerEngines BladeEngine2 10Gbps network devices has been updated to version 4.0.100u.
- The ixgbevf driver has been updated to version 2.1.0-k
- The cxgb4 driver for Chelsio Terminator4 10G Unified Wire Network Controllers has been updated.
- The cxgb3 driver for the Chelsio T3 Family of network devices has been updated.
- The ixgbe driver for Intel 10 Gigabit PCI Express network devices is updated to version 3.4.8-k.
- The e1000e driver for Intel PRO/1000 network devices has been updated to version 1.3.16-k.
- The e1000 driver for Intel PRO/1000 network devices has been updated, providing support for Marvell Alaska M88E1118R PHY.
- The e100 driver has been updated.
- The enic driver for Cisco 10G Ethernet devices has been updated to version 2.1.1.24.
- The igbvf driver has been updated to version 2.0.0-k.
- The igb driver for Intel Gigabit Ethernet Adapters has been updated.
- The bnx2 driver for the NetXtreme II 1 Gigabit Ethernet controllers has been updated to version 2.1.6+.
- The tg3 driver for Broadcom Tigon3 ethernet devices has been updated to version 3.119.
- The qlcnic driver for the HP NC-Series QLogic 10 Gigabit Server Adapters has been updated to version 5.0.16+.
- The bna driver has been updated.
- The r8169 driver has been updated to fix two bugs related to Rx checksum offloading.
- The qlge driver has been updated to version v1.00.00.29.
- The cnic driver has been updated to add iSCSI and FCoE support for the 578xx family of Multi-Port Single-Chip 10G Ethernet Converged Controllers, VLAN support, and new bnx2x firmware interface.

-
- The `iwl6000` and `iwlwifi` have been update with the EEPROM version 0x423.

Graphics and miscellaneous drivers

- The `radeon` driver has been updated with post-3.0 fixes, including backported `drm/agp` code.
- The `nouveau` and `i915` drivers have been updated, including backported `drm/agp` code.
- The Ricoh memory stick driver (`R5C592`) has been update with new `KFIFO` application programming interface.
- The `netjet` driver has been updated to blacklist the Digium TDM400P PCI Card.
- The `lm78` driver has been updated.
- The `wacom` driver has been updater to add support for the Cintiq 21UX2, Intuos4 WL, and DTU-2231 adapter cards.
- The `synaptics` driver has been updated to add multi-touch support.
- The ALSA HDA audio driver has been updated to enable or improve support for new chipsets and HDA audio codecs.
- The `edac` driver has been updated to support the new Northbridge chip for AMD platforms.

Storage

iprutil support for SAS VRAID functions

The *iprutils* package provides utilities to manage and configure SCSI devices that are supported by the *ipr* SCSI storage device driver. The *iprutils* package has been updated to support SAS VRAID functions for new 6 GB SAS adapters on IBM POWER7.

LVM RAID support

In Red Hat Enterprise Linux 6.2, support for MD's RAID personalities has been added to LVM as a *Technology Preview*. The following basic features are available: create, display, rename, use, and remove RAID logical volumes. Automated fault tolerance is not yet available.

It is possible to create RAID logical volumes by specifying the `--type <segtype>` argument. The following are a few examples:

- Create a RAID1 array (this is a different implementation of RAID1 than LVM's `mirror` segment type):

```
~]# lvcreate --type raid1 -m 1 -L 1G -n my_lv my_vg
```

- Create a RAID5 array (3 stripes + 1 implicit parity):

```
~]# lvcreate --type raid5 -i 3 -L 1G -n my_lv my_vg
```

- Create a RAID6 array (3 stripes + 2 implicit parity):

```
~]# lvcreate --type raid6 -i 3 -L 1G -n my_lv my_vg
```

iSCSI extension for RDMA (iSER) initiator and target

iSER initiator and target is now fully supported. Red Hat Enterprise Linux can now function as an iSCSI initiator and storage server in production environments that use InfiniBand and where high throughput and low latency are key requirements.

Reduced activation times for LVM devices

LVM devices can now be activated or deactivated quicker than before. This is relevant to high-density environments that involve a large number of LVM configurations. An example of this is a host that supports hundreds of virtual guests each using one or more logical volumes.

File System

XFS scalability

The XFS file system is currently supported in Red Hat Enterprise Linux 6 and is well suited for very large files and file systems on a single host. Integrated backup and restore, direct I/O and online resizing of the file system are some of the benefits that this file system provides.

The XFS implementation has been improved to better handle metadata intensive workloads. An example of this type of workload is accessing thousands of small files in a directory. Prior to this enhancement, metadata processing could cause a bottleneck and lead to degraded performance. To address this problem an option to delay the logging of the metadata has been added that provides a significant performance improvement. As a result of this delayed logging of metadata, XFS performance is on par with ext4 for such workloads. The default mount options have also been updated to use delayed logging.

Parallel NFS

Parallel NFS (pNFS) is a part of the NFS v4.1 standard that allows clients to access storage devices directly and in parallel. The pNFS architecture eliminates the scalability and performance issues associated with NFS servers in deployment today.

pNFS supports 3 different storage protocols or layouts: files, objects and blocks. The Red Hat Enterprise Linux 6.2 NFS client supports the files layout protocol.

To automatically enable the pNFS functionality, create the `/etc/modprobe.d/dist-nfsv41.conf` file with the following line and reboot the system:

```
alias nfs-layouttype4-1 nfs_layout_nfsv41_files
```

Now when the `-o minorversion=1` mount option is specified, and the server is pNFS-enabled, the pNFS client code is automatically enabled.

This feature is a Technology Preview. For more information on pNFS, refer to <http://www.pnfs.com/>.

Asynchronous writes in CIFS

The CIFS (Common Internet File System) protocol allows for a unified way to accessing remote files on disparate operating systems. The CIFS client has traditionally only allowed for synchronous writes. This meant that the client process would not yield back control until the writes were successfully completed. This can lead to degraded performance for large transactions that take long to complete. The CIFS client has been updated to write data in parallel without the need to wait for the sequential writes. This change can now result in performance improvements up to 200%.

CIFS NTLMSSP authentication

Support for NTLMSSP authentication has been added to CIFS. In addition, CIFS now uses the kernel's crypto API.

autofs4 module

The `autofs4` module has been updated to kernel version 2.6.38.

Fixed tracepoints for ext3 and jbd

Fixed tracepoints have been added to ext3 and jbd.

Mount options in superblock

Support for the **-o nobarrrier** mount option in ext4, and its utilities: **tune2fs**, **debugfs**, **libext2fs**, has been added.

Networking

Multi-message send system call

Red Hat Enterprise Linux 6.2 introduces the multi-message **send** system call which is the **send** version of the existing **recvmmsg** system call in Red Hat Enterprise Linux 6.

The system call **sendmmsg** socket API looks like this:

```
struct mmsghdr {
    struct msghdr msg_hdr;
    unsigned msg_len;
};

ssize_t sendmmsg(int socket, struct mmsghdr *datagrams, int vlen, int flags);
```

Transmit Packet Steering (XPS)

Red Hat Enterprise Linux 6.2 includes Transmit Packet Steering (XPS) for multiqueue devices. XPS introduces more efficient transmission of network packets for multiqueue devices by specifically targeting the processor involved in sending the packet. XPS enables the selection of the transmit queue for packet transmission based on configuration. This is analogous to the receive-side functionality implemented in Red Hat Enterprise Linux 6.1 which allowed for processor selection based on the receive queue (RPS). XPS has shown to improve throughput by 20% to 30%.

Traffic flooding for unregistered groups

Previously, the bridge flooded packets to unregistered groups to all ports. However, this behavior is not desirable in environments where traffic to unregistered groups is always present. In Red Hat Enterprise Linux 6.2, traffic is only sent to unregistered groups to ports marked as routers. To force flooding to any given port, mark that port as a router.

Stream Control Transmission Protocol (SCTP) Multihome support

Red Hat Enterprise Linux 6.2 adds support for SCTP multihoming—the ability of nodes (that is, multi-home nodes) to be reached under several IP addresses.

Tracepoints for UDP packet drop events

In Red Hat Enterprise Linux 6.2, more tracepoints have been added for UDP packet drop events. These tracepoints provide a way to analyze the reasons why UDP packets are dropped.

IPSet

The IPSet feature in the kernel has been added to store multiple IP addresses or port numbers, and match them against a collection via **iptables**.

TCP initial receive window default

The TCP initial receive window default has been increased from 4 kB to 15 kB. The benefit of this increase is that any data (15 kB > payload > 4 kB) can now fit in the initial window. With a 4 kB setting (IW3), any payload larger than 4 kB would have to be broken into multiple transfers.

TCP initial congestion window default

In Red Hat Enterprise Linux 6.2, the TCP initial congestion window default is now set to 10, according to [RFC 5681](http://tools.ietf.org/html/rfc5681)¹. Additionally, the initial-window code common to TCP and CCID-2 has been consolidated.

GSO support on IPv6

GSO (Generic Segmentation Offload) support for the IPv6 forward path has been added, improving the performance of host/guest communication if GSO is enabled.

vios-proxy

vios-proxy is a stream-socket proxy for providing connectivity between a client on a virtual guest and a server on a Hypervisor host. Communication occurs over virtio-serial links. This feature is introduced as a Technology Preview in Red Hat Enterprise Linux 6.2.

¹ <http://tools.ietf.org/html/rfc5681>

Authentication and Interoperability

Identity Management

Red Hat Enterprise Linux 6.2 includes identity management capabilities that allow for central management of user identities, policy-based access control and authentication services. This identity management service, previously referred to as IPA, is based on the open source FreeIPA project. These services have been present as a Technology Preview in previous Red Hat Enterprise Linux 6 releases. With this release, identity management has been promoted to fully supported.



Further Reading

The *Identity Management Guide*¹ provides detailed information about the Identity Management solution, the technologies with which it works, and some of the terminology used to describe it. It also provides high-level design information for both the client and server components.

PIV support for smart cards

Support for smart cards with a PIV (Personal Identity Verification) interface has been added in Red Hat Enterprise Linux 6.2. It is now possible to use FIPS 201 compliant PIV cards that allow for secure use of data. PIV cards enable confidentiality of data by restricting access to the card holder. They also ensure data integrity by allowing only the card holder to make modifications. They guarantee the authenticity of the information and prevent non-repudiation of data. The use of PIV cards is mandated via the U.S. Homeland Security Presidential Directive 12 (HSPC-12) which requires the use of this type of technology to gain access to all government IT systems.

¹ https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Identity_Management_Guide/index.html

Entitlement

Certificate-based RHN

The new subscription management platform delivers Red Hat subscriptions and software services in a flexible, scalable and secure way. When installing a new Red Hat Enterprise Linux 6 system, a user receives X.509 certificates which contain information about which Red Hat products are installed and what subscriptions the machine is consuming. The subscription information includes support levels, expiration dates, Red Hat account numbers, and Red Hat contract numbers. In addition, an X.509 certificate allows a machine to authenticate to the Red Hat Content Delivery Network (CDN). The globally distributed Red Hat Content Delivery Network (CDN) is designed to work even with an outage of the Red Hat systems. Users outside of North America should see improved update speeds and availability with the new system. RHN Classic continues to be the default option for computer registration and for receiving updates.

Entitlement certificates for disconnected systems

Red Hat customer portal, in a conjunction with new functionality available in Red Hat Enterprise Linux 6.2, allows customers to register and subscribe up to 25 machines that are completely disconnected. Prior to this enhancement, customers with disconnected systems were not able to receive benefits from subscription information and tracking from the RHN website. For customers with more than 25 disconnected machines, RHN Satellite continues to be a recommended option at additional cost.

Automatic regeneration of a certificate after renewal of a subscription

It is now possible to automatically regenerate new entitlement certificates after the renewal of a subscription. Prior to this enhancement, customers were required to manually regenerate the certificate to continue to receive software updates and other subscription services. Automatically regenerating a certificate minimizes service interruptions. Users are also notified of cases where auto regeneration of certificates was not successful. For more information, refer to <https://www.redhat.com/rhel/renew/faqs/>.

Red Hat Subscription Manager and the Subscription Service

In Red Hat Enterprise Linux 6.2, during system registration, the Red Hat Subscription Manager is now used by default.



Further Reading

The [Red Hat Enterprise Linux 6.2 Deployment Guide](https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/entitlements.html)¹ contains further information on managing subscriptions.

The [Red Hat Enterprise Linux 6.2 Installation Guide](https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Installation_Guide/sn-firstboot-updates.html)² contains further information on the registration and subscription process during **firstboot** and **kickstart**.

¹ https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/entitlements.html

² https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Installation_Guide/sn-firstboot-updates.html

Security, Standards and Certification

Common Criteria Certification

As of Red Hat Enterprise Linux 6.2 Beta, Red Hat Enterprise Linux 6 is under evaluation for Common Criteria at Evaluation Assurance Level (EAL) 4+. Common Criteria provides a standardized way of expressing security requirements and defines a set of rigorous criteria by which products are evaluated.

FIPS-140 Validation

As of Red Hat Enterprise Linux 6.2 Beta, Red Hat Enterprise Linux 6 cryptographic modules are under evaluation for FIPS-140 certification. FIPS-140 is a U.S. government security standard used for accrediting cryptographic modules. Red Hat Enterprise Linux now satisfies the regulatory requirement mandated by the U.S. federal government for acceptable use of cryptographic modules by all the governmental agencies.

Trusted boot

Red Hat Enterprise Linux 6.2 includes Intel Trusted Boot, a trusted boot mechanism (provided by the *tboot* package). Trusted boot is an install-time optional component that allows for Intel's Trusted Execution Technology (TXT) to perform a measured and verified launch of the operating system kernel. Trusted boot is supported on both Intel x86 and Intel 64 architectures.

Compiler and Tools

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 6.2 is updated to version 1.6, providing:

- Kernel modules with a hyphen ("-") in their name, such as `i2c-core` are now handled properly.
- `process.mark` now supports `$$parms` for reading probe parameters.
- Improved and simplified operation of the SystemTap **compile-server** and **client**:
 - **compile-server** may cache script build results for improved performance.
 - **compile-server** and **client** communicate exchange version information to adjust the communication protocol accordingly and use the newest version of the server possible.
 - Removal of deprecated tools: **stap-client**, **stap-authorize-server-cert**, **stap-authorize-signing-cert**, **stap-find-or-start-server**, and **stap-find-servers**.
- For remote execution, the `--remote USER@HOST` functionality can now be specified multiple times and will automatically build the script for distinct kernel and architecture configurations, and run it on all named machines at once.
- The **staprun** utility now allows multiple instances of the same script to be run at same time.

Clustering

Dynamic schema generation

The introduction of dynamic schema generation provides a lot of flexibility for end users to plug into Red Hat Enterprise Linux High Availability Add-on custom resource and fence agents, and still retain the possibility to validate their `/etc/cluster.conf` configuration file against those agents. It is a strict requirement that custom agents provide correct metadata output and that the agents must be installed on all cluster nodes.

Clustered Samba on GFS2

Support for Samba in a clustered environment is now fully supported in Red Hat Enterprise Linux 6.2. Samba clustering relies on a clustered file system being available and shared on all nodes. In the Red Hat Enterprise Linux context, Samba clustering has been configured to work with GFS2, the native shared storage file system.

Clustered Samba (more specifically CTDB) provides the ability for the metadata to span multiple physical hosts in a cluster. CTDB will automatically recover and repair node-specific databases in case of node failures. It also provides high availability features like node monitoring and failover.

Support for redundant ring for standalone Corosync

Red Hat Enterprise Linux 6.2 introduces support for redundant ring with autorecovery feature as a Technology Preview. Refer to the [Technical Notes](#)¹ for a list of known issues associated with this Technology Preview.

corosync-cpgtool

The `corosync-cpgtool` now specifies both interfaces in a dual ring configuration. This feature is a Technology Preview.

Disabling rgmanager in /etc/cluster.conf

As a consequence of converting the `/etc/cluster.conf` configuration file to be used by `pacemaker`, `rgmanager` must be disabled. The risk of not doing this is high; after a successful conversion, it would be possible to start `rgmanager` and `pacemaker` on the same host, managing the same resources.

Consequently, Red Hat Enterprise Linux 6.2 includes a feature (as a Technology Preview) that forces the following requirements:

- `rgmanager` must refuse to start if it sees the `<rm disabled="1">` flag in `/etc/cluster.conf`.
- `rgmanager` must stop any resources and exit if the `<rm disabled="1">` flag appears in `/etc/cluster.conf` during a reconfiguration.

¹ https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/6.2_Technical_Notes/index.html

High Availability

XFS on High Availability Add On

Usage of XFS in conjunction with Red Hat Enterprise Linux 6.2 High Availability Add On as a file system resource is now fully supported.

HA support for VMWare

Applications running inside VMWare based guests can now be configured for high availability. This also includes full support for the use of GFS2 shared storage file system in the environment. A new SOAP-based fence agent has been added that has the ability to fence guests when necessary.

Administrative UI enhancements

Luci, the web-based administrative UI for configuring clusters has been updated to include the following:

- Role-based access control (RBAC): enables fine-grained access levels by defining user classes to access specific cluster operations.
- Improved response times for destructive operations in a cluster.

Support for UDP-Unicast

IP multicasting has been the only supported option for a cluster transport. IP multicasting is inherently complex to configure and often requires re-configuration of network switches. UDP-unicast in contrast offers a simpler approach to cluster configuration and is an established protocol for cluster communication. UDP-unicast, initially introduced as a Technology Preview, is now fully supported.

Watchdog integration with fence_scsi

Watchdog is a general timer service available in Linux that can be used to periodically monitor system resources. Fence agents have now been integrated with watchdog such that the watchdog service can reboot a node after it has been fenced using `fence_scsi`. This eliminates the need for manual intervention to reboot the node after it has been fenced using `fence_scsi`.

Virtualization

KVM processor performance improvement

Virtual CPU timeslice sharing

Virtual CPU timeslice sharing is a performance enhancing feature at the Linux scheduler level, where a spinning virtual CPU can hand the remainder of its timeslice to another virtual CPU before yielding the CPU. This feature addresses an inherent lock holder preemption issue that exists in SMP systems, that can affect performance in virtual CPUs. This feature provides a stable performance in multi-processor guests. This feature is supported on both Intel and AMD processors, and is called Pause Loop Exiting (PLE) on Intel processors, and Pause Filter on AMD processors.

KVM network performance improvements

KVM network performance is a critical requirement for Virtualization and cloud based products and solutions. Red Hat Enterprise Linux 6.2 provides a number of network performance optimizations to improve the KVM network para-virtualized driver performance in various setups.

Improved small message KVM performance

Red Hat Enterprise Linux 6.2 improves the KVM small message performance to satisfy a variety of networking workloads that generate small messages (< 4K).

Wire speed requirement in KVM network drivers

Virtualization and cloud products that run networking work loads need to run wire speeds. Up until Red Hat Enterprise Linux 6.1, the only way to reach wire speed on a 10 GB Ethernet NIC with a lower CPU utilization was to use PCI device assignment (passthrough), which limits other features like memory overcommit and guest migration

The **macvtap/vhost** zero-copy capabilities allows the user to use those features when high performance is required. This feature improves performance for any Red Hat Enterprise Linux 6.x guest in the VEPA use case. This feature is introduced as a Technology Preview.

UDP checksum optimization for KVM network drivers

UDP checksum optimization eliminates the need for the guest to validate the checksum if it has been validated by host NICs. This feature speeds up UDP on external to guest on 10 GB Ethernet cards with Red Hat Enterprise Linux 6.2 guests and hosts. The UDP checksum optimization is implemented in the `virtio-net` driver.

Improved I/O path performance when host slower than guest

The Red Hat Enterprise Linux 6.2 KVM network driver has improved I/O path performance, with reduced virtual machine exits and interrupts, that results in faster data delivery. This also improvement enables you to run a faster guest on a slower host, without incurring any performance penalties. This enhancement is achieved by an enhanced `virtio` ring structure, and event index support in `virtio` and `vhost-net`.

KVM Systems Management and usability improvements

System monitoring via SNMP

This feature provides KVM support for stable technology that is already used in data center with bare metal systems. SNMP is the standard for monitoring and is extremely well understood as well as computationally efficient. System monitoring via SNMP in Red Hat Enterprise Linux 6.2 allows the KVM hosts to send SNMP traps on events so that hypervisor events can be communicated

to the user via standard SNMP protocol. This feature is provided through the addition of a new package: *libvirt-snmpp*. This feature is introduced as a Technology Preview.

Improved guest debugging capabilities

Users who virtualize their data centers need a way of debugging when a guest OS hangs up and a crash dump has to be initiated. There are two methods heavily used with physical systems:

- Triggering a non-maskable interrupt (NMI) in the guest
- Sending SysRq sequences to the guest

While these capabilities are provided directly with the KVM console, a number of users use KVM through the **libvirt** API and **virsh**, where these two features were missing. Red Hat Enterprise Linux 6.2 improves guest debugging capabilities across the KVM stack, thus allowing a user to trigger NMIs in guests and send SysRq key sequences to guests.

Improve virtual machine boot up access

Users who virtualize their data centers need to track the guest boot up process and display the entire BIOS and kernel boot up message from the start. The absence of this feature prevents users from an interactive use of the **virsh** console, prior to boot up. A new package, *sgabios*, has been added to Red Hat Enterprise Linux 6.2, to provide this capability, along with some additions to **qemu-kvm**.

Live Snapshots

Red Hat Enterprise Linux 6.2 introduces the Live Snapshot feature as a Technology Preview. The live snapshots feature provides automatic backup of virtual machine images on the hard drive, and provides a per drive snapshot of the virtual disks transparently, using external *qcow2* images. Multi-disk live snapshot creation helps with maintaining data integrity by pausing *qemu* prior to taking as many snapshots as there are disks. Thus, a multi-disk snapshot will have all disks containing data from the same point in time.

It is important to know that there is a limitation with file system consistency. However, the re-use of the snapshot image is crash-consistent. A user would have to run a file system check (**fsck**) or replay journal entries, which is similar to booting after pulling the power cord.

Multi-processor (NUMA) Tuning Improvements

Red Hat Enterprise Linux 6.2 adds tuning improvements to the **libvirt** API stack, resulting in improved out-of-the-box performance when performing SPECvirt measurements. Red Hat Enterprise Linux 6.2 is now able to pin the memory associated with a NUMA node when a virtual machine is created.

USB enhancements

The USB 2.0 emulation has been implemented for **qemu-kvm**. This is available for QEMU directly only. Libvirt support is planned for the next release.

Remote Wakeup support has been added for the USB host controller. Together with the cooperation of the guest OS it allows to stop the frequent 1000hz polling mode and put the device into sleep. It dramatically improves the power utilization and the CPU consumption of virtual machines with a USB mouse emulation (or a tablet) — one of the common devices that every virtual machine has.

Xen improvements

Memory ballooning

Memory ballooning is now supported by Red Hat Enterprise Linux 6 paravirtualized Xen guests.

Domain memory limit

Memory limit for x86_64 domU PV guests has been increased to 128 GB:
CONFIG_XEN_MAX_DOMAIN_MEMORY=128.

Time accounting

The `xen_sched_clock` implementation (which returns the number of unstolen nanoseconds) has been replaced by the `xen_clocksource_read` implementation.

Virtualization Documentation

The Red Hat Enterprise Linux Virtualization Guide has been divided into several specific guides:

- [Red Hat Enterprise Linux Virtualization Getting Started Guide](#)¹
- [Red Hat Enterprise Linux Virtualization Administration Guide](#)²
- [Red Hat Enterprise Linux Virtualization Host Configuration and Guest Installation Guide](#)³

spice-protocol

The package `spice-protocol` has been upgraded to version 0.8.1, providing the following new features:

- Support for volume change
- Support for async guest I/O writes and interrupts
- Support for suspend (S3) related guest I/O writes
- Support for interrupt indicating a guest bug

Linux Containers

Linux containers provide a flexible approach to application runtime containment on bare-metal systems without the need to fully virtualize the workload. Red Hat Enterprise Linux 6.2 provides application level containers to separate and control the application resource usage policies via cgroup and namespaces. This release introduces basic management of container life-cycle by allowing creation, editing and deletion of containers via the **libvirt** API and the **virt-manager** GUI. Linux Containers are a Technology Preview.

Red Hat Enterprise Virtualization Hypervisor RPM multi-installable

In order to allow side-by-side installs of the `rhev-hypervisor` package, configure Yum to make `rhev-hypervisor` an installonly package by editing the `/etc/yum.conf` file and adding the `installonlypkgs` option:

```
[main]
...
installonlypkgs=rhev-hypervisor
```

This option needs to also include the default list of installonly packages which can be found in the `yum.conf` man page (`man yum.conf 5`) in the `installonlypkgs` option section.

Graphics

The X server shipped in Red Hat Enterprise Linux 6.2 has been updated to the upstream X.org 1.10 X server and the upstream Mesa 7.11 releases. The X server had internal structure changes that required updating all video and input drivers. In addition, the kernel graphics support has been updated to include new hardware support and bug fixes.

AMD

Improved support for ATI/AMD GPU series HD2xxx, HD4xxx, HD5xxx, FirePro. Support added for new HD6xxx series, new model in the FirePro series and new mobile GPU HD6xxxM series.

Intel

Support has been added for Intel's next generation class of chipsets.

Nouveau

2D/Xv acceleration is now supported on GeForce GT2xx (and Quadro equivalents). Suspend/resume support has been improved.

X server

RandR-enabled drivers (intel, nouveau, radeon) now confine the cursor to the visible area of the screen in asymmetric multihead configurations.

The Composite extension is now functional when Xinerama is used to span a single desktop across multiple GPUs.

X server configuration may now be managed with configuration file snippets under `/etc/X11/xorg.conf.d/` in addition to `/etc/X11/xorg.conf` itself. X.org input device configuration in these snippets applies when the device becomes available to the X server at runtime.

Refer to the 1.10 X server upstream announcement for additional information: <http://lists.freedesktop.org/archives/xorg-announce/2011-February/001612.html>.

Mesa

See the Mesa 7.11 upstream announcement for release notes: <http://mesa3d.org/relnotes-7.11.html>.

General Updates

Matahari

Matahari in Red Hat Enterprise Linux 6.2 is fully supported only for x86 and AMD64 architectures. Builds for other architectures are considered a Technology Preview.

Automatic Bug Reporting Tool

Red Hat Enterprise Linux 6.2 introduces **ABRT** 2.0. **ABRT** logs details of software crashes on a local system, and provides interfaces (both graphical and command line based) to report issues to various issue trackers, including Red Hat support. This update provides the following notable enhancements:

- More flexible configuration with a new syntax.
- Out-of-process plugins (plugins run in separate processes and communicate via inter-process communication with other processes). Advantages of such a design are:
 - bugs in plugins do not break the main daemon,
 - more secure as most of the processing is now done under the normal (non-root) user,
 - plugins can be written in any programming language.
- Reporting backend is shared across all of Red Hat's issue reporting tools:
 - **ABRT**, **sealert**, all users of **python-meh** (**Anaconda**, **firstboot**)
 - Because all of the tools above share the same configuration, it only has to be written once.



Further Reading

For more information on ABRT configuration and its new syntax, refer to the [Red Hat Enterprise Linux 6.2 Deployment Guide](https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/ch-abrt.html)¹.

Optimized math library for Linux on IBM System z

Red Hat Enterprise Linux 6.2 provides an optimized linear algebra math library for Linux on System z which enables the compiler to generate code for high profile functions, taking advantage of the latest hardware functions.

Improved tablet support

Red Hat Enterprise Linux 6.2 improves support for Wacom devices. It is no longer necessary to reconfigure device settings after a device has been unplugged and plugged back in.

Improved wireless detection

NetworkManager can now scan wireless networks in the background, providing a better user experience.

¹ https://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/ch-abrt.html

Increase in CPU support in GNOME

The `gnome-system-monitor` utility can now monitor systems that have more than 64 CPUs.

Appendix A. Component Versions

This appendix is a list of components and their versions in the Red Hat Enterprise Linux 6.2 release.

Table A.1. Component Versions

Component	Version
Kernel	2.6.32-202
QLogic qla2xxx driver	8.03.07.05.06.2-k
QLogic qla2xxx firmware	ql23xx-firmware-3.03.27-3.1 ql2100-firmware-1.19.38-3.1 ql2200-firmware-2.02.08-3.1 ql2400-firmware-5.06.01-1 ql2500-firmware-5.06.01-1
Emulex lpfc driver	8.3.5.45.2p
iSCSI initiator utils	6.2.0.872-27
DM-Multipath	0.4.9-43
LVM	2.02.87-3
X Server	1.10.4-3

Appendix B. Revision History

Revision 1-0 Tue Dec 6 2011

Martin Prpič mpropic@redhat.com

Release of the Red Hat Enterprise Linux 6.2 Release Notes

