



# Red Hat Enterprise Linux 6

## 6.3 Release Notes

Release Notes for Red Hat Enterprise Linux 6.3

Edition 3

Last Updated: 2017-10-20



# Red Hat Enterprise Linux 6 6.3 Release Notes

---

Release Notes for Red Hat Enterprise Linux 6.3  
Edition 3

Red Hat Engineering Content Services

## Legal Notice

Copyright © 2012 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 6.3. For detailed documentation on all changes to Red Hat Enterprise Linux for the 6.3 update, refer to the Technical Notes.

---

## Table of Contents

<b>PREFACE</b> .....	<b>2</b>
<b>CHAPTER 1. KERNEL</b> .....	<b>3</b>
<b>CHAPTER 2. DEVICE DRIVERS</b> .....	<b>6</b>
<b>CHAPTER 3. NETWORKING</b> .....	<b>10</b>
<b>CHAPTER 4. RESOURCE MANAGEMENT</b> .....	<b>11</b>
<b>CHAPTER 5. AUTHENTICATION AND INTEROPERABILITY</b> .....	<b>12</b>
<b>CHAPTER 6. SUBSCRIPTION MANAGEMENT</b> .....	<b>14</b>
<b>CHAPTER 7. VIRTUALIZATION</b> .....	<b>15</b>
7.1. KVM .....	15
7.2. SPICE .....	17
7.3. LIBVIRT .....	17
<b>CHAPTER 8. CLUSTERING AND HIGH AVAILABILITY</b> .....	<b>18</b>
<b>CHAPTER 9. STORAGE</b> .....	<b>19</b>
<b>CHAPTER 10. GENERAL UPDATES</b> .....	<b>21</b>
<b>APPENDIX A. COMPONENT VERSIONS</b> .....	<b>23</b>
<b>APPENDIX B. REVISION HISTORY</b> .....	<b>24</b>

## PREFACE

Red Hat Enterprise Linux minor releases are an aggregation of individual enhancement, security and bug fix errata. The Red Hat Enterprise Linux 6.3 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 changes (that is, bugs fixed, enhancements added, and known issues found) in this minor release are available in the [Technical Notes](#). The Technical Notes document also contains a complete list of all currently available Technology Previews along with packages that provide them.



### IMPORTANT

The online *Red Hat Enterprise Linux 6.3 Release Notes* which are located online [here](#), are to be considered the definitive, up-to-date version. Customers with questions about the release are advised to consult the online *Release* and *Technical Notes* for their version of Red Hat Enterprise Linux.

Should you require information regarding the Red Hat Enterprise Linux life cycle, refer to <https://access.redhat.com/support/policy/updates/errata/>.

# CHAPTER 1. KERNEL

The kernel shipped in Red Hat Enterprise Linux 6.3 includes several hundred bug fixes for, and enhancements to, the Linux kernel. For details concerning every bug fixed and every enhancement added to the kernel for this release, refer to the kernel section of the [Red Hat Enterprise Linux 6.3 Technical Notes](#).

## Thin-provisioning and scalable snapshot capabilities

The `dm-thin` targets, `thin` and `thin-pool`, provide a device-mapper device with thin-provisioning and scalable snapshot capabilities. This feature is available as a Technology Preview. For more information on the newly-introduced LVM thin provisioning, refer to [Chapter 9, Storage](#).

## sysfs mbox interface deprecated

The `lpfc` driver is deprecating the `sysfs mbox` interface as it is no longer used by the Emulex tools. Read and write operations are now stubbed out and only return the `-EPERM` (Operation not permitted) symbol.

## Supported Kdump targets

For a complete list of supported Kdump targets (that is, targets that kdump can use to dump a vmcore to), refer to the following Kbase article: <https://access.redhat.com/knowledge/articles/41534>.

## Support for additional mount options

Red Hat Enterprise Linux 6.3 adds support for mount options to restrict access to `/proc/<PID>/` directories. One of the new options is called `hidepid=` and its value defines how much information about processes is provided to non-owners. The `gid=` option defines a group that gathers information about all processes. Untrusted users, which are not supposed to monitor tasks in the whole system, should not be added to the group.

## O\_DIRECT flag support

Support for the `O_DIRECT` flag for files in FUSE (File system in Userspace) has been added. This flag minimizes cache effects of the I/O to and from a file. In general, using this flag degrades performance, but it is useful in special situations, such as when applications do their own caching.

## CONFIG\_STRICT\_DEVMEM enabled on PowerPC

In Red Hat Enterprise Linux 6.3, the `CONFIG_STRICT_DEVMEM` configuration option is enabled by default for the PowerPC architecture. This option restricts access to the `/dev/mem` device. If this option is disabled, userspace access to all memory is allowed, including kernel and userspace memory, and accidental memory (write) access could potentially be harmful.

## CONFIG\_HPET\_MMAP enabled

In Red Hat Enterprise Linux 6.3, the high-resolution timer's capacity to remap the HPET registers into the memory of a user process has been enabled.

## Improved performance on large systems

A number of patches have been applied to the kernel in Red Hat Enterprise Linux 6.3 to improve overall performance and reduce boot time on extremely large systems (patches were tested on a system with 2048 cores and 16 TB of memory).

## rdrand kernel support

The Intel Core i5 and i7 processors (formerly code-named Ivy Bridge) support a new `rdrand` instruction to quickly generate random numbers. The kernel shipped in Red Hat Enterprise Linux 6.3 utilizes this instruction to provide quick random number generation.

### UEFI support for persistent storage

Persistent storage (`pstore`), a file system interface for platform dependent persistent storage, now supports UEFI.

### CPU family specific container files

Support for CPU family specific container files has been added. Starting with AMD family 15h processors, a container such as `microcode_amd_fam15h.bin` is now loaded for the aforementioned family of processors.

### USB 3.0 support

Red Hat Enterprise Linux 6.3 includes full USB 3.0 support.

### kdump/kexec kernel dumping mechanism for IBM System z

In Red Hat Enterprise Linux 6.3, the `kdump/kexec` kernel dumping mechanism is enabled for IBM System z systems as a Technology Preview, in addition to the IBM System z stand-alone and hypervisor dumping mechanism. The auto-reserve threshold is set at 4 GB; therefore, any IBM System z system with more than 4 GB of memory has the `kexec/kdump` mechanism enabled.

Sufficient memory must be available because `kdump` reserves approximately 128 MB as default. This is especially important when performing an upgrade to Red Hat Enterprise Linux 6.3. Sufficient disk space must also be available for storing the dump in case of a system crash. `Kdump` is limited to DASD or QETH networks as dump devices until `kdump` on SCSI disk is supported.

The following warning message may appear when `kdump` is initialized:

```
..no such file or directory
```

This message does not impact the dump functionality and can be ignored. You can configure or disable `kdump` via `/etc/kdump.conf`, `system-config-kdump`, or `firstboot`.

### Module-accessible interface for ftrace

The `ftrace` function tracer now allows modules and all users to make use of the `ftrace` function tracing utility. For more information, refer to the following man pages:

```
man trace-cmd-record
man trace-cmd-stack
```

### Tracing of multi-threaded processes

When tracing processes with more than one thread, the `ltrace` utility would neglect to trace threads other than the main thread. But because threads share address space, those other threads would still see breakpoints distributed by `ltrace`. Consequently, those threads would be killed by a `SIGTRAP` signal. Red Hat Enterprise Linux 6.3 includes thread-awareness and breakpoint handling mechanisms. Support for tracing of multi-threaded processes is now on par with tracing single-threaded process.

### Cross Memory Attach

Cross Memory Attach provides a mechanism to reduce the number of data copies needed for intra-node inter-process communication. In particular, this allows MPI libraries doing intra-node communication to do a single copy of the message rather than a double copy of the message via shared



memory. This technique has been employed in the past through multiple, unique driver-based implementations. The implementation introduced in Red Hat Enterprise Linux 6.3 provides a general solution for this functionality. In addition, it provides a layer of abstraction for device driver writers who wish to exploit these functions without having to modify their corresponding implementations when there are changes in the memory management subsystem.

### Spinning mutex performance enhancement for IBM System z

Red Hat Enterprise Linux 6.3 enhances the usage of mutexes. Additional information provided to the scheduler allows for more efficient and less costly decisions when optimizing processor cycles depending on the usage of mutexes, thread scheduling, and the status of the physical and virtual processors. The status of a thread owning a locked mutex is examined and waiting threads are not scheduled unless the first thread is scheduled on both a virtual and a physical processor.

### Enhanced DASD statistics for PAV and HPF for IBM System z

Red Hat Enterprise Linux 6.3 enables improved diagnosis of PAV (Parallel Access Volume) and HPF (High Performance Ficon) environments to analyze and tune the DASD performance on a system, for example, to give recommendations regarding the number of alias devices or the usage of PAV versus HyperPAV.

### OSA concurrent software/hardware trap for IBM System z

With Red Hat Enterprise Linux 6.3, collective problem analysis through consolidated dumps of software and hardware is enabled. A command can be used to generate `qeth` or `qdio` trace data and to trigger the internal dump of an OSA device.

### Added ability to switch between two graphics cards

The `CONFIG_VGA_SWITCHEROO` configuration option is now enabled by default to allow switching between two graphics cards.

### KEXEC\_AUTO\_THRESHOLD lowered to 2 GB

With Red Hat Enterprise Linux 6.3, the `crashkernel=auto` parameter changed the default kdump enabling threshold from 4 GB to 2 GB. This means that any machine that has 2 GB (or more) of memory will have the kdump feature enabled on its systems.

Available memory in the system is rounded up to 128 MB when calculating the 2 GB threshold for deciding whether to enable kdump or not. If a system has 1920 MB (2G-128M) RAM available, kdump is enabled.

Run the following commands if you wish to disable kdump (for example, due to memory constraints):

1. To stop the `kdump` service, execute the following command:

```
~]# service kdump stop
```

2. To disable the `kdump` service, execute the following command:

```
~]# chkconfig kdump off
```

3. To return memory previously reserved for `kdump` back to the system, execute the following command:

```
~]# echo 0 > /sys/kernel/kexec_crash_size
```

## CHAPTER 2. DEVICE DRIVERS

### BFA driver fully supported

The Brocade BFA Fibre Channel and FCoE driver is no longer a Technology Preview. In Red Hat Enterprise Linux 6.3 the BFA driver is fully supported.

### BNA driver fully supported

The Brocade BNA driver for Brocade 10Gb PCIe ethernet Controllers is no longer a Technology Preview. In Red Hat Enterprise Linux 6.3 the BNA driver is fully supported.

### SR-IOV on the be2net driver

The SR-IOV functionality of the Emulex **be2net** driver is considered a Technology Preview in Red Hat Enterprise Linux 6.3. You must meet the following requirements to use the latest version of SR-IOV support:

- You must run the latest Emulex firmware (revision 4.1.417.0 or later).
- The server system BIOS must support the SR-IOV functionality and have virtualization support for Direct I/O VT-d.
- You must use the GA version of Red Hat Enterprise Linux 6.3.

SR-IOV runs on all Emulex-branded and OEM variants of BE3-based hardware, which all require the **be2net** driver software.

### Storage drivers

- Red Hat Enterprise Linux 6.3 includes the **mtip32xx** driver which adds support for Micron RealSSD P320h PCIe SSD drives.
- The **lpfc** driver for Emulex Fibre Channel Host Bus Adapters has been updated to version 8.3.5.68.2p.
- The **mptfusion** driver has been updated to version 3.04.20.
- The **bnx2fc** for the Broadcom Netxtreme II 57712 chip has been updated to version 1.0.11.
- The **qla2xxx** driver for QLogic Fibre Channel HBAs has been updated to version 8.04.00.02.06.3-k. The **qla2xxx** driver update for Red Hat Enterprise Linux 6.3 now takes advantage of the common code in the SCSI mid-layer that handles queue-full status messages returned from a target port. Before, this code resided in the **qla2xxx** driver itself. To maintain API compatibility, stubs for the **ql2xqfulltracking** and **ql2xqfullrampup** module parameters have been left in the driver itself.

In addition, this update also adds support for ISP82xx and ISP83xx, and adds the dynamic logging functionality.

- The **qla4xxx** has been updated to version 5.02.00.00.06.03-k1, which adds support for displaying **port\_state**, **port\_speed**, and **targetalias** in the **sysfs** file system.
- The **megaraid** driver has been updated to version 00.00.06.14-rh1.
- The **ipr** driver for IBM Power Linux RAID SCSI HBAs has been updated to enable the SAS VRAID functions.

- The `cciss` driver has been updated to add older controllers to the `kdump` blacklist.
- The `hpsa` driver has been updated to the version 2.0.2-4 to add older controllers to the `kdump` blacklist.
- The `bnx2i` driver for Broadcom NetXtreme II iSCSI has been updated to version 2.7.2.1.
- The `mpt2sas` driver has been updated to version 12.101.00.00, adding NUNA I/O support which uses multi-reply queue support of HBAs.
- The `mptsas` driver has been updated to add the following device ID: `SAS1068_820XELP`.
- The Brocade BFA FC SCSI driver (`bfa` driver) has been updated.
- The `be2iscsi` driver for ServerEngines BladeEngine 2 Open iSCSI devices has been updated.
- The `ahci.c` driver have been updated to add the AHCI-mode SATA DeviceID for the Intel DH89xxCC PCH.
- The `iscsi` driver has been updated to version 1.1 to pick up the latest Intel hardware support, enhancements, and bug fixes.
- The `iscsi_sata` driver has been updated to add T10 DIF support.
- The `libfc`, `libfcOE`, and `fcOE` drivers have been updated to fix various bugs and add several enhancements.
- The `libsas` driver has been updated.
- The `qib` driver for TrueScale HCAs has been updated.
- The `libata` module has been updated to fix various bugs.
- The `dm-raid` code of the `md` driver has been updated to include flush support.
- The following drivers have been updated to the latest version: `ahci`, `md/bitmap`, `raid0`, `raid1`, `raid10`, and `raid456`.
- The `aacraid` driver has been updated to version 1.1-7[28000].

### Network drivers

- The `netxen` driver for NetXen Multi port (1/10) Gigabit Network has been updated to version 4.0.77 or greater.
- The `bnx2x` driver has been updated to version 7.2.16 to include support for the 578xx family of chips.
- The `be2net` driver for ServerEngines BladeEngine2 10Gbps network devices has been updated to version 4.2.5.0r.
- The `ixgbev` driver has been updated to version 2.2.0-k to include the latest hardware support, enhancements, and bug fixes.
- 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 has been updated to version 3.6.7-k to include the latest hardware support, enhancements, and bug fixes.
- The **e1000e** driver for Intel PRO/1000 network devices has been updated.
- The **e1000** driver for Intel PRO/1000 network devices has been updated.
- The **e100** driver has been updated.
- The **enic** driver for Cisco 10G Ethernet devices has been updated to version 2.1.1.35, adding SR-IOV support.
- The **igbvf** driver (Intel Gigabit Virtual Function Network driver) has been updated to version 2.0.1-k.
- The **igb** driver for Intel Gigabit Ethernet Adapters has been updated to version 3.2.10-k, providing up-to-date hardware support, enhancements, and bug fixes.
- The **bnx2** driver for the NetXtreme II 1 Gigabit Ethernet controllers has been updated to version 1.0.11.
- The **tg3** driver for Broadcom Tigon3 Ethernet devices has been updated to version 3.120+.
- The **qlcnic** driver for the HP NC-Series QLogic 10 Gigabit Server Adapters has been updated to version 5.0.26.
- The **bnx** driver has been updated.
- The **r8169** driver has been updated to add support for the latest Realtek NICs (8168D/8168DP/8168E/8168EV) and increase reliability of older NICs.
- The **qlge** driver has been updated to version 1.00.00.30.
- The **cnic** driver has been updated to version 2.5.9, which improves error recovery on bnx2x devices, adds FCoE parity error recovery, increases the maximum amount of FCoE sessions, and adds other enhancements.
- The **iwl6000** and **iwlwifi** drivers have been updated to add support for the Intel Centrino Wireless-N 6235 series of Wi-Fi adapters. The **iwlwifi** also adds an option with 5GHz band can be disabled.
- The wireless LAN subsystem has been updated. It introduces the **dma\_unmap** state API and adds a new kernel header file: **include/linux/pci-dma.h**.
- The **at11c** driver has been updated to the latest upstream version, which adds support for Atheros AR8151 v2 and Atheros AR8152 PCI-E Gigabit Ethernet Controllers.

#### Miscellaneous drivers

- The **i915** driver has been updated.
- Various graphics drivers have been updated with DRM support rebased to version 3.3-rc2.

- The `wacom` driver has been updated, deprecating the `wacompl` package and obsoleting the `wdaemon` package.
- The ALSA HDA audio driver has been updated to enable or improve support for new chipsets and HDA audio codecs.
- The `btusb` driver has been updated to include support for the Broadcom BCM20702A0 single-chip bluetooth processor.
- The `k10temp` driver from the `hwmon` subsystem has been updated to add support for AMD family 12h/14h/15h of CPUs.
- The ALPS Touchpad driver has been updated to add support for ALPS Touchpad protocol versions 3 and 4, and add support for touchpads with 4-directional buttons.
- The `jsm` driver has been updated to add Enhanced Error Handling (EEH).
- The `m1x4_en` driver has been updated to version 2.0.
- The `m1x4_core` driver has been updated to version 1.1.

## CHAPTER 3. NETWORKING

### QFQ queuing discipline

In Red Hat Enterprise Linux 6.3, the `tc` utility has been updated to work with the Quick Fair Scheduler (QFQ) kernel features. Users can now take advantage of the new QFQ traffic queuing discipline from userspace. This feature is considered a Technology Preview.

### rdma\_bw and rdma\_last utilities deprecated

The `rdma_bw` and `rdma_lat` utilities (provided by the `perftest` package) are now deprecated and will be removed from the `perftest` package in a future update. Users should use the following utilities instead: `ib_write_bw`, `ib_write_lat`, `ib_read_bw`, and `ib_read_lat`.

### Enhancements in the configuration tool for IBM System z network devices

With Red Hat Enterprise Linux 6.3, the System z `qethconf` tool provides information messages when a change of attributes did not work as expected.

### IPv6 support for qetharp tool for IBM System z

Red Hat Enterprise Linux 6.3 adds IPv6 support to the `qetharp` tool for inspection and modification of the ARP cache of HiperSockets (real and virtual) operated in layer 3 mode. For real HiperSockets, the tool queries and shows the IPv6 address, and for guest LAN HiperSockets, it queries and shows IPv6 to MAC address mappings.

### Maximum number of session slots

In Red Hat Enterprise Linux 6.3, NFSv4 introduces a module/kernel boot parameter, `nfs.max_session_slots`, which sets the maximum number of session slots the NFS client will attempt to negotiate with the server. This limits the number of simultaneous RPC requests that the client can send to the NFSv4 server. Note that setting this value higher than `max_tcp_slot_table_limit` has no effect.

## CHAPTER 4. RESOURCE MANAGEMENT

### Network priority cgroup resource controller

Red Hat Enterprise Linux 6.3 introduces the Network Priority (`net_prio`) resource controller, which provides a way to dynamically set the priority of network traffic per each network interface for applications within various cgroups. For more information, refer to the *Resource Management Guide*

### OOM control and notification API for cgroups

The memory resource controller implements an Out-of-Memory (OOM) notifier which uses the new notification API. When enabled (by executing `echo 1 > memory.oom_control`), an application is notified via `eventfd` when an OOM occurs. Note that OOM notification does not function for root cgroups.

### New numad package

The numad package provides a daemon for NUMA (Non-Uniform Memory Architecture) systems that monitors NUMA characteristics. As an alternative to manual static CPU pinning and memory assignment, numad provides dynamic adjustment to minimize memory latency on an ongoing basis. The package also provides an interface that can be used to query the numad daemon for the best manual placement of an application. The numad package is introduced as a Technology Preview.

## CHAPTER 5. AUTHENTICATION AND INTEROPERABILITY

### Support for central management of SSH keys

Previously, it was not possible to centrally manage host and user SSH public keys. Red Hat Enterprise Linux 6.3 includes SSH public key management for Identity Management servers as a Technology Preview. OpenSSH on Identity Management clients is automatically configured to use public keys which are stored on the Identity Management server. SSH host and user identities can now be managed centrally in Identity Management.

### SELinux user mapping

Red Hat Enterprise Linux 6.3 introduces the ability to control the SELinux context of a user on a remote system. SELinux user map rules can be defined and, optionally, associated with HBAC rules. These maps define the context a user receives depending on the host they are logging into and the group membership. When a user logs into a remote host which is configured to use SSSD with the Identity Management backend, the user's SELinux context is automatically set according to mapping rules defined for that user. For more information, refer to [http://freeipa.org/page/SELinux\\_user\\_mapping](http://freeipa.org/page/SELinux_user_mapping). This feature is considered a Technology Preview.

### Multiple required methods of authentication for sshd

SSH can now be set up to require multiple ways of authentication (whereas previously SSH allowed multiple ways of authentication of which only one was required for a successful login); for example, logging in to an SSH-enabled machine *requires* both a passphrase and a public key to be entered. The `RequiredAuthentications1` and `RequiredAuthentications2` options can be configured in the `/etc/ssh/sshd_config` file to specify authentications that are required for a successful log in. For example:

```
~]# echo "RequiredAuthentications2 publickey,password" >>
/etc/ssh/sshd_config
```

For more information on the aforementioned `/etc/ssh/sshd_config` options, refer to the `sshd_config` man page.

### SSSD support for automount map caching

In Red Hat Enterprise Linux 6.3, SSSD includes a new Technology Preview feature: support for caching automount maps. This feature provides several advantages to environments that operate with `autofs`:

- Cached automount maps make it easy for a client machine to perform mount operations even when the LDAP server is unreachable, but the NFS server remains reachable.
- When the `autofs` daemon is configured to look up automount maps via SSSD, only a single file has to be configured: `/etc/sss/sss.conf`. Previously, the `/etc/sysconfig/autofs` file had to be configured to fetch `autofs` data.
- Caching the automount maps results in faster performance on the client and lower traffic on the LDAP server.

### Change in SSSD `debug_level` behavior

SSSD has changed the behavior of the `debug_level` option in the `/etc/sss/sss.conf` file. Previously, it was possible to set the `debug_level` option in the `[sss]` configuration section and the result would be that this became the default setting for other configuration sections, unless they explicitly overrode it.



Several changes to internal debug logging features necessitated that the `debug_level` option must always be specified independently in each section of the configuration file, instead of acquiring its default from the `[sssd]` section.

As a result, after updating to the latest version of SSSD, users may need to update their configurations in order to continue receiving debug logging at the same level. Users that configure SSSD on a per-machine basis can use a simple Python utility that updates their existing configuration in a compatible way. This can be accomplished by running the following command as root:

```
~]# python /usr/lib/python2.6/site-packages/sss_update_debug_levels.py
```

This utility makes the following changes to the configuration file: it checks to see if the `debug_level` option was specified in the `[sssd]` section. If so, it adds that same level value to each other section in the `sss.conf` file for which `debug_level` is unspecified. If the `debug_level` option already exists explicitly in another section, it is left unchanged.

Users who rely on central configuration management tools need to make these same changes manually in the appropriate tool.

#### New `ldap_chpass_update_last_change` option

A new option, `ldap_chpass_update_last_change`, has been added to SSSD configuration. If this option is enabled, SSSD attempts to change the *shadowLastChange* LDAP attribute to the current time. Note that this is only related to a case when the LDAP password policy is used (usually taken care of by LDAP server), that is, the LDAP extended operation is used to change the password. Also note that the attribute has to be writable by the user who is changing the password.

## CHAPTER 6. SUBSCRIPTION MANAGEMENT

### Migration from RHN Classic to certificate-based RHN

Red Hat Enterprise Linux 6.3 includes a new tool to migrate RHN Classic customers to the certificate-based RHN. For more information, refer to the [Red Hat Enterprise Linux 6 Subscription Management Guide](#).

### Subscription Manager `gpgcheck` behavior

Subscription Manager now disables `gpgcheck` for any repositories it manages which have an empty `gpgkey`. To re-enable the repository, upload the GPG keys, and ensure that the correct URL is added to your custom content definition.

### Firstboot System Registration

In Red Hat Enterprise Linux 6.3, during `firstboot` system registration, registering to Certificate-based Subscription Management is now the default option.

### Server side deletes

System profiles are now unregistered when they are deleted from the Customer Portal so that they no longer check in with certificate-based RHN.

### Preferred service levels

Subscription manager now allows users to associate a machine with a preferred *Service Level* which impacts the auto subscription and healing logic. For more information on service levels, refer to the [Red Hat Enterprise Linux 6 Subscription Management Guide](#) .

### Limiting updates to a specific minor release

Subscription manager now allows a user to select a specific release (for example, Red Hat Enterprise Linux 6.2), which will lock a machine to that release. Prior to this update, there was no way to limit package updates in the event newer packages became available as part of a later minor release (for example, Red Hat Enterprise Linux 6.3).

# CHAPTER 7. VIRTUALIZATION

## 7.1. KVM

### KVM scalability enhancements

KVM scalability enhancements in Red Hat Enterprise Linux 6.3 include:

- The maximum supported virtual guest size increased from 64 to 160 virtual CPUs (vCPUs).
- The maximum supported memory in a KVM guest increased from 512 GB to 2 TB.

### KVM support for new Intel and AMD processors

KVM in Red Hat Enterprise Linux 6.3 includes support for:

- Intel Core i3, i5, i7 and other processors formerly code named “Sandy Bridge”,
- and new AMD family 15h processors (code named “Bulldozer”).

The new CPU model definitions in KVM provide the necessary new processor enablement to KVM host and the virtualized guests. This ensures that KVM Virtualization derives the performance benefits associated with the new processors, and takes advantage of the new instructions in the latest CPUs.

### KVM “Steal Time” support

*Steal time* is the time that a virtual CPU waits for a real CPU while the hypervisor is servicing another virtual processor. KVM virtual machines can now calculate and report steal time, visible through tools like `top` and `vmstat`, which provides a guest with accurate CPU utilization data.

The KVM steal time feature provides accurate data to a guest regarding CPU utilization and virtual machine performance. Large amounts of steal time indicates that the virtual machine performance is curtailed by the CPU time assigned to the guest by the hypervisor. The user can relieve the performance issues caused by CPU contention by running fewer guests on the host or by increasing the CPU priority of the guest. The KVM steal time value provides users with data to allow them to take the next step in improving their application run-time performance.

### Improved access to qcow2 disk images

KVM in Red Hat Enterprise Linux 6.3 improved the access to `qcow2` disk images (`qcow2` is the default format) by making it more asynchronous, thus avoiding vCPU stalls and enhancing the overall performance during disk I/O.

### New qemu-guest-agent subpackage

The `qemu-guest-agent` package can be installed on virtual guest systems to provide the `qemu-ga` service. The `qemu-ga` service starts automatically (launching the `/usr/bin/qemu-ga` daemon) if the `/dev/virtio-ports/org.qemu.guest_agent.0` file exists. The daemon can be used to respond to a variety of requests for information and actions on the guests and it is currently encapsulated by `libvirt` on Red Hat Enterprise Linux 6 systems.

The `qemu-ga` daemon is used by `libvirt` to request that the guest VM suspend to disk or suspend to RAM. In addition to suspend operations, the daemon can respond to shutdown commands and file system freeze requests during a virtual machine live snapshot (to get a consistent disk state).

### Performance monitoring in KVM guests

KVM can now virtualize Intel's performance monitoring unit (PMU) to allow virtual machines to use performance monitoring.

Note that the `-cpu host` flag must be set when using this feature.

With this feature, Red Hat virtualization customers running Red Hat Enterprise Linux 6 guests can use the CPU's PMU counter while using the performance tool for profiling. The virtual performance monitoring unit feature allows virtual machine users to identify sources of performance problems in their guests, thereby improving the ability to profile a KVM guest from the host.

This feature is a Technology Preview in Red Hat Enterprise Linux 6.3 and is only supported with guests running Red Hat Enterprise Linux 6.

### **Dynamic virtual CPU allocation**

KVM in Red Hat Enterprise Linux 6.3 now supports dynamic virtual CPU allocation, also called vCPU hot plug, to dynamically manage capacity and react to unexpected load increases on their platforms during off-peak hours.

The virtual CPU hot-plugging feature gives system administrators the ability to dynamically adjust CPU resources in a guest. Because a guest no longer has to be taken offline to adjust the CPU resources, the availability of the guest is increased.

This feature is a Technology Preview in Red Hat Enterprise Linux 6.3. Currently, only the vCPU hot-add functionality works. The vCPU hot unplug feature is not yet implemented.

### **Virtio-SCSI capabilities**

KVM Virtualization's storage stack has been improved with the addition of virtio-SCSI (a storage architecture for KVM based on SCSI) capabilities. Virtio-SCSI provides the ability to connect directly to SCSI LUNs and significantly improves scalability compared to virtio-blk. The advantage of virtio-SCSI is that it is capable of handling hundreds of devices compared to virtio-blk which can only handle 28 devices and exhausts PCI slots.

Virtio-SCSI is now capable of inheriting the feature set of the target device with the ability to:

- attach a virtual hard drive or CD through the virtio-scsi controller,
- pass-through a physical SCSI device from the host to the guest via the QEMU scsi-block device,
- and allow the usage of hundreds of devices per guest; an improvement from the 28-device limit of virtio-blk.

This feature is a Technology Preview in Red Hat Enterprise Linux 6.3.

### **Support for in-guest S4/S3 states**

KVM's power management features have been extended to include native support for S4 (suspend to disk) and S3 (suspend to RAM) states within the virtual machine, speeding up guest restoration from one of these low power states. In earlier implementations guests were saved or restored to or from a disk or memory that was external to the guest, which introduced latency.

Additionally, guests can be awakened from the S3 state with events from a remote keyboard through SPICE.

This feature is a Technology Preview and is disabled by default in Red Hat Enterprise Linux 6.3. To enable it, select the `/usr/share/seabios/bios-pm.bin` file for the VM bios instead of the default `/usr/share/seabios/bios.bin` file.

The native, in-guest S4 (suspend to disk) and S3 (suspend to RAM) power management features support the ability to perform suspend to disk and suspend to RAM functions in the guest (as opposed

to the host), reducing the time needed to restore a guest by responding to simple keyboard input. This also removes the need to maintain an external memory-state file. This capability is supported on Red Hat Enterprise Linux 6.3 guests and Windows guests running on any hypervisor capable of supporting S3 and S4.

### SR-IOV support for NIC

Red Hat Enterprise Linux 6.3 introduces SR-IOV support for network interface controllers. This feature allows a NIC on a KVM host to be shared by KVM guests. For more information on SR-IOV, refer to *Chapter 13. SR-IOV* in the [Virtualization Host Configuration and Guest Installation Guide](#). For information on SR-IOV on the `be2net` driver, refer to *Chapter 2, Device Drivers*.

### TSC scaling in KVM for AMD-V

Red Hat Enterprise Linux 6.3 adds support for Time Stamp Counter (TSC) scaling to KVM for AMD Virtualization (AMD-V). This feature is capable of emulating a given TSC frequency on a KVM guest.

### Support for perf-kvm

Support for the `perf-kvm` tool, which provides the ability to monitor guest performance from host, has been added. For more information, refer to the `perf-kvm` man page.

## 7.2. SPICE

### USB 2.0 redirection support

Spice builds on KVM USB 2.0 host adapter emulation support, and enables remote USB redirection support that allows virtual machines running on servers to use remotely plugged USB devices on the client side.

## 7.3. LIBVIRT

### Controlling up/down link states

`libvirt` is now capable of controlling the state (up or down) of a link of the guest virtual network interfaces. This allows users to perform testing and simulation as though plugging and unplugging the network cable from the interface. This feature also lets users isolate guests in case any problems arise.

### Added support for the latest Intel and AMD processors

In Red Hat Enterprise Linux 6.3, `libvirt` has been updated to add support for the latest Intel Core i3, i5, i7 and other Intel processors, and family 15h microarchitecture AMD processors. With this update, `libvirt` now utilizes the new features these processors include.

## CHAPTER 8. CLUSTERING AND HIGH AVAILABILITY

### Administrative UI enhancements

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

- A confirmation dialog box appears when removing a clustered service.
- The UI includes an improved restart icon.
- The **Add a child resource** button has been simplified.
- An option to enable debugging from the UI has been added.

### Automatic timeout of inactive luci authenticated sessions

As of Red Hat Enterprise Linux 6.3, **luci** authenticated sessions automatically time out after 15 minutes of inactivity. This period can be configured in the `/etc/sysconfig/luci` file by modifying the `who.auth_tkt_timeout` parameter.

### New libqb package

The **libqb** package provides a library with the primary purpose of providing high performance client server reusable features, such as high performance logging, tracing, inter-process communication, and polling. This package is introduced as a dependency of the **pacemaker** package, and is considered a Technology Preview in Red Hat Enterprise Linux 6.3.

### Pacemaker now uses libqb logging

Because of the newly added **libqb** dependency, **pacemaker** now uses its logging functionality to provide less verbosity while keeping the ability to debug and support **pacemaker**.

### Utilizing CPG API for inter-node locking

**Rgmanager** includes a feature which enables it to utilize Corosync's Closed Process Group (CPG) API for inter-node locking. This feature is automatically enabled when Corosync's Redundant Ring Protocol (RRP) feature is enabled. Corosync's RRP feature is considered fully supported. However, when used with the rest of the High-Availability Add-Ons, it is considered a Technology Preview.

## CHAPTER 9. STORAGE

### LVM support for (non-clustered) thinly-provisioned snapshots

A new implementation of LVM copy-on-write (cow) snapshots is available in Red Hat Enterprise Linux 6.3 as a Technology Preview. The main advantage of this implementation, compared to the previous implementation of snapshots, is that it allows many virtual devices to be stored on the same data volume. This implementation also provides support for arbitrary depth of recursive snapshots (snapshots of snapshots of snapshots ...).

This feature is for use on a single-system. It is not available for multi-system access in cluster environments.

For more information, refer to documentation of the `-s/--snapshot` option in the `lvcreate` man page.

### LVM support for (non-clustered) thinly-provisioned LVs

Logical Volumes (LVs) can now be thinly provisioned to manage a storage pool of free space to be allocated to an arbitrary number of devices when needed by applications. This allows creation of devices that can be bound to a thinly provisioned pool for late allocation when an application actually writes to the LV. The thinly-provisioned pool can be expanded dynamically if and when needed for cost-effective allocation of storage space. In Red Hat Enterprise Linux 6.3, this feature is introduced as a Technology Preview. You must have the `device-mapper-persistent-data` package installed to try out this feature. For more information, refer to the `lvcreate` man page.

### Dynamic aggregation of LVM metadata via `lvmetad`

Most LVM commands require an accurate view of the LVM metadata stored on the disk devices on the system. With the current LVM design, if this information is not available, LVM must scan all the physical disk devices in the system. This requires a significant amount of I/O operations in systems that have a large number of disks.

The purpose of the `lvmetad` daemon is to eliminate the need for this scanning by dynamically aggregating metadata information each time the status of a device changes. These events are signaled to `lvmetad` by `udev` rules. If `lvmetad` is not running, LVM performs a scan as it normally would.

This feature is provided as a Technology Preview and is disabled by default in Red Hat Enterprise Linux 6.3. To enable it, refer to the `use_lvmetad` parameter in the `/etc/lvm/lvm.conf` file, and enable the `lvmetad` daemon by configuring the `lvm2-lvmetad` init script.

### Fiber Channel over Ethernet (FCoE) target mode fully supported

Fiber Channel over Ethernet (FCoE) target mode is fully supported in Red Hat Enterprise Linux 6.3. This kernel feature is configurable via the `targetcli` utility, 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 `targetcli(8)` man pages (provided by the `lldpad` and `fcoe-target-utils` packages, respectively).

### LVM RAID fully supported with the exception of RAID logical volumes in HA-LVM

The expanded RAID support in LVM is now fully supported in Red Hat Enterprise Linux 6.3. LVM is now capable of creating RAID 4/5/6 logical volumes and supports mirroring of these logical volumes. The MD (software RAID) modules provide the backend support for these new features.

### Activating volumes in read-only mode

A new LVM configuration file parameter, *activation/read\_only\_volume\_list*, makes it possible to always activate particular volumes in read-only mode, regardless of the actual permissions on the volumes concerned. This parameter overrides the `--permission rw` option stored in the metadata.



## CHAPTER 10. GENERAL UPDATES

### Matahari packages deprecated

The Matahari agent framework (matahari-\*) packages are deprecated starting with the Red Hat Enterprise Linux 6.3 release. Focus for remote systems management has shifted towards the use of the CIM infrastructure. This infrastructure relies on an already existing standard which provides a greater degree of interoperability for all users. It is strongly recommended that users discontinue the use of the matahari packages and other packages which depend on the Matahari infrastructure (specifically, libvirt-qmf and fence-virt-d-libvirt-qpid). It is recommended that users uninstall Matahari from their systems to remove any possibility of security issues being exposed.

Users who choose to continue to use the Matahari agents should note the following:

- The matahari packages are not installed by default starting with Red Hat Enterprise Linux 6.3 and are not enabled by default to start on boot when they are installed. Manual action is needed to both install and enable the `matahari` services.
- The default configuration for `qpid` (the transport agent used by Matahari) does not enable access control lists (ACLs) or SSL. Without ACLs/SSL, the Matahari infrastructure is not secure. Configuring Matahari without ACLs/SSL is not recommended and may reduce your system's security.
- The `matahari-services` agent is specifically designed to allow remote manipulation of services (start, stop). Granting a user access to Matahari services is equivalent to providing a remote user with root access. Using Matahari agents should be treated as equivalent to providing remote root SSH access to a host.
- By default in Red Hat Enterprise Linux, the Matahari broker (`qpidd` running on port 49000) does not require authentication. However, the Matahari broker is not remotely accessible unless the firewall is disabled, or a rule is added to make it accessible. Given the capabilities exposed by Matahari agents, if Matahari is enabled, system administrators should be extremely cautious with the options that affect remote access to Matahari.

Note that Matahari will not be shipped in future releases of Red Hat Enterprise Linux (including Red Hat Enterprise Linux 7), and may be considered for formal removal in a future release of Red Hat Enterprise Linux 6.

### Software Collections utilities

Red Hat Enterprise Linux 6.3 includes an `scl-utils` package which provides a runtime utility and packaging macros for packaging Software Collections. Software Collections allow users to concurrently install multiple versions of the same RPM packages on the system. Using the `scl` utility, users may enable specific versions of RPMs which are installed in the `/opt` directory. For more information on Software Collections, refer to the *Software Collections Guide*.

### The `openssl-ibmca` package is now part of the IBM System z default installation

With Red Hat Enterprise Linux 6.3, the `openssl-ibmca` package is part of the System z default installation. This avoids the need for manual installation steps.

### MySQL InnoDB plug-in

Red Hat Enterprise Linux 6.3 provides the MySQL InnoDB storage engine as a plug-in for AMD64 and Intel 64 architectures. The plugin offers additional features and better performance than the built-in InnoDB storage engine.

### OpenJDK 7

Red Hat Enterprise Linux 6.3 includes full support for OpenJDK 7 as an alternative to OpenJDK 6. The `java-1.7.0-openjdk` packages provide the OpenJDK 7 Java Runtime Environment and the OpenJDK 7 Java Software Development Kit.

### New Java 7 packages

The `java-1.7.0-oracle` and `java-1.7.0-ibm` packages are now available in Red Hat Enterprise Linux 6.3.

### Setting the NIS domain name via initscripts

The `initscripts` package has been updated to allow users to set the NIS domain name. This is done by configuring the `NISDOMAIN` parameter in the `/etc/sysconfig/network` file, or other relevant configuration files.

### ACL support for logrotate

Previously, when certain groups were permitted to access all logs via ACLs, these ACLs were removed when the logs were rotated. In Red Hat Enterprise Linux 6.3, the `logrotate` utility supports ACLs, and logs that are rotated preserve any ACL settings.

### The `wacomcpl` package deprecated

The `wacomcpl` package has been deprecated and has been removed from the package set. The `wacomcpl` package provided graphical configuration of Wacom tablet settings. This functionality is now integrated into the GNOME Control Center.

### Updated NumPy package

The NumPy package which is designed to manipulate large multi-dimensional arrays of arbitrary records has been updated to version 1.4.1. This updated version includes these changes:

- When operating on `0-d` arrays, `numpy.max` and other functions accept only the following parameters: `axis=0`, `axis=-1`, and `axis=None`. Using out-of-bounds axes indicates a bug, for which NumPy now raises an error.
- Specifying the `axis > MAX_DIMS` parameter is no longer allowed; NumPy now raises an error, instead of behaving the same as when `axis=None` was specified.

### Rsyslog updated to major version 5

The `rsyslog` package has been upgraded to major version 5. This upgrade introduces various enhancements and fixes multiple bugs. The following are the most important changes:

- The `$HUPisRestart` directive has been removed and is no longer supported. Restart-type HUP processing is therefore no longer available. Now, when the SIGHUP signal is received, outputs (log files in most cases) are only re-opened to support log rotation.
- The format of the spool files (for example, disk-assisted queues) has changed. In order to switch to the new format, drain the spool files, for example, by shutting down `rsyslogd`. Then, proceed with the Rsyslog upgrade, and start `rsyslogd` again. Once upgraded, the new format is automatically used.
- When the `rsyslogd` daemon was running in the debug mode (using the `-d` option), it ran in the foreground. This has been fixed and the daemon is now forked and runs in the background, as is expected.

For more information on changes introduced in this version of Rsyslog, refer to <http://www.rsyslog.com/doc/v5compatibility.html>.

## APPENDIX A. COMPONENT VERSIONS

This appendix is a list of components and their versions in the Red Hat Enterprise Linux 6.3 release.

**Table A.1. Component Versions**

Component	Version
Kernel	2.6.32-279
QLogic qla2xxx driver	8.04.00.02.06.3-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-5.06.05-1 ql2500-firmware-5.06.05-1
Emulex lpfc driver	8.3.5.68.2p
iSCSI initiator utils	6.2.0.872-41
DM-Multipath	0.4.9-56
LVM	2.02.95-10

## APPENDIX B. REVISION HISTORY

<b>Revision 6-0.2</b>	<b>Mon Feb 18 2013</b>	<b>Martin Prpič</b>
Removed incorrect information about <code>rt2800usb/rt2x00</code> being updated.		
<b>Revision 1-0</b>	<b>Wed Jun 20 2012</b>	<b>Martin Prpič</b>
Release of the Red Hat Enterprise Linux 6.3 Release Notes.		
<b>Revision 0-0</b>	<b>Tue Apr 24 2012</b>	<b>Martin Prpič</b>
Release of the Red Hat Enterprise Linux 6.3 Beta Release Notes.		