Red Hat Enterprise Linux 6

6.4 Release Notes

Release Notes for Red Hat Enterprise Linux 6.4

Edition 4

Red Hat Enterprise Linux 6 6.4 Release Notes

Release Notes for Red Hat Enterprise Linux 6.4
Edition 4

Red Hat Engineering Content Services

Legal Notice

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 6.4. For detailed documentation on all changes to Red Hat Enterprise Linux for the 6.4 update, refer to the Technical Notes.

Table of Contents

PREFACE

Red Hat Enterprise Linux minor releases are an aggregation of individual enhancement, security and bug fix errata. The *Red Hat Enterprise Linux 6.4 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 6.4 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. INSTALLATION

FCoE Support in the Kickstart File

When using a kickstart file to install Red Hat Enterprise Linux 6.4, with the new **_fcoe_** kickstart option you can specify which Fibre Channel over Ethernet (FCoE) devices should be activated automatically in addition to those discovered by Enhanced Disk Drive (EDD) services. For more information, refer to the *Kickstart Options* section in the Red Hat Enterprise Linux 6 *Installation Guide*.

Installation over VLAN

In Red Hat Enterprise Linux 6.4, the `vlanid=` boot option and the `--vlanid=` kickstart option allow you to set a virtual LAN ID (802.1q tag) for a specified network device. By specifying either one of these options, installation of the system can be done over a VLAN.

Configuring Bonding

The **bond** boot option and the `--bondslaves` and `--bondopts` kickstart options can now be used to configure bonding as a part of the installation process. For more information on how to configure bonding, refer to the following parts of the Red Hat Enterprise Linux 6 *Installation Guide*: section *Kickstart Options* and chapter *Boot Options*.

CHAPTER 2. KERNEL

The kernel shipped in Red Hat Enterprise Linux 6.4 includes several hundred bug fixes for, and enhancements to, the Linux kernel. For details concerning important bugs fixed and enhancements added to the kernel for this release, refer to the kernel section of the Red Hat Enterprise Linux 6.4 Technical Notes.

Fibre Channel Protocol: End-To-End Data Consistency Checking

Data integrity between a host adapter and a storage server has been improved in Red Hat Enterprise Linux 6.4 by implementing the zFCP-specific part of the enhanced T10 DIF SCSI standard for End-To-End (E2E) data consistency checking.

Flash Express Support for IBM System z

Storage-Class Memory (SCM) for IBM System z is a class of data storage devices that combine properties of both storage and memory. SCM for System z now supports Flash Express memory. SCM increments can be accessed through Extended Asynchronous Data Mover (EADM) subchannels. Each increment is represented by a block device. This feature improves the paging rate and access performance for temporary storage, for example for data warehousing.

Open vSwitch Kernel Module

Red Hat Enterprise Linux 6.4 includes the Open vSwitch kernel module as an enabler for Red Hat's layered product offerings. Open vSwitch is supported only in conjunction with those products containing the accompanying user space utilities. Please note that without these required user space utilities, Open vSwitch will not function and can not be enabled for use. For more information, please refer to the following Knowledge Base article: https://access.redhat.com/knowledge/articles/270223.

Oracle ASMLib Availability and Support

Oracle ASM (Automated Storage Management) is a data volume manager for Oracle databases. ASMLib is an optional utility that can be used on Linux systems to manage Oracle ASM devices. ASMLib consists of the following components:

- kmod-oracleasm (open-source (GPL) kernel module package)

- oracleasm-support (open-source (GPL) utilities package)

- oracleasmlib (proprietary library package)

ASM features and functionality are available without ASMLib. The use of ASMLib does not affect database performance. The ASMLib kernel module package is now available in the Red Hat Enterprise Linux 6 Supplementary RHN channel. Red Hat Enterprise Linux 6 customers who use ASMLib can obtain the other two components using the instructions in the following KnowledgeBase article:

https://access.redhat.com/knowledge/solutions/315643

The ASMLib kernel module package is provided for the convenience of our customers via the "Supplementary" Red Hat Network (RHN) channel. Red Hat's support team will field ASMLib related calls and use commercially reasonable efforts to support the ASMLib kernel module until such efforts require knowledge of or modifications to Oracle's proprietary dependent component(s). That said, the ASMLib kernel module package ABIs are not guaranteed per Red Hat's Supplementary software package support terms available at:

https://access.redhat.com/support/offerings/production/scope_moredetail.html

Please note that Red Hat is continuing to develop fully open-source alternatives to ASMLib. Red Hat has provided a reference architecture for Oracle RAC clusters using upstream-accepted technologies such as dm-multipath and udev. This reference architecture is available at:

https://access.redhat.com/knowledge/articles/216093

Comparison of Booted System and Dumped System

This feature allows you to compare a booted system with a dumped system to efficiently analyze changes that might be introduced by image migration. To identify a guest, `stsi` and `stfle` data is used. A new function, `lgr_info_log()` compares the current data (`lgr_info_cur`) with the last recorded one (`lgr_info_last`).

Perf Tool Updated

The **perf** tool has been updated to upstream version 3.6-rc7, which provides a large number of bug fixes and enhancements. The following is a list of notable enhancements:

- Kprobe events support was added.

- A new **perf** event command line syntax engine has been included, which allows curly brackets (`{` and `}`) to be used for definition of event groups, for example: `{cycles,cache-misses}`.

- The **perf** annotate browser has been enhanced to allow navigation through ASM calls and jumps.

- The **perf** tool has been updated to provide a per-user view with the new `--uid` command line option. When used, **perf** shows tasks for a specified user only.

- The **perf** tool now provides a wider variety of automated tests.

Uncore PMU Support

The kernel shipped with Red Hat Enterprise Linux 6.4 adds "uncore" Performance Monitoring Unit (PMU) support to the perf event subsystem for Intel Xeon Processor X55xx and Intel Xeon Processor X56xx family of processors. The "uncore" refers to subsystems in the physical processor package that are shared by multiple processor cores, for example the L3 cache. With uncore PMU support, performance data can be easily collected on a package level.

PMU events parsing has also been enabled to allow debugging via perf.

Reduced `memcg` Memory Overhead

Memory control groups maintain their own Least Recently Used (LRU) list to, for example, reclaim memory. This list was on top of the global per-zone LRU list. In Red Hat Enterprise Linux 6.4, the memory overhead for `memcg` was reduced by disabling the global per-zone LRU list and converting its users to operate on the per-memory cgroup lists instead.

Memory Reclaim and Compaction

The kernel shipped with Red Hat Enterprise Linux 6.4 uses reclaim and compaction for high-order allocation requests or under memory pressure.

Support of the Transactional Execution Facility and Runtime Instrumentation Facility

Support of the Transactional-Execution Facility (available with IBM zEnterprise EC12) in the Linux kernel helps eliminate software locking overhead that can impact performance and offer increased scalability and parallelism to drive higher transaction throughput. Support of the Runtime Instrumentation Facility (available with IBM zEnterprise EC12) provides an advanced mechanism to profile program code for improved analysis and optimization of the code generated by the new IBM JVM.

Fail-open Mode

Red Hat Enterprise Linux 6.4 adds support for a new fail-open mode when using netfilter's NFQUEUE target. This mode allows users to temporarily disable packet inspection and maintain connectivity under heavy network traffic.

kdump and kexec Kernel Dumping Mechanism for IBM System z Fully Supported

In Red Hat Enterprise Linux 6.4, the kdump/kexec kernel dumping mechanism is enabled for IBM System z systems as a fully supported feature, 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 kdump/kexec mechanism enabled.

Sufficient memory must be available because kdump reserves approximately 128 MB by default. This is especially important when performing an upgrade to Red Hat Enterprise Linux 6.4. Sufficient disk space must also be available for storing the dump in case of a system crash.

You can configure or disable kdump through **/etc/kdump.conf**, **system-config-kdump**, or **firstboot**.

TSC Deadline Support for KVM

TSC deadline timer is a new mode in the Local APIC (LAPIC) timer, which generates one-shot timer interrupts based on the TSC deadline, in place of the current APIC clock count interval. It provides more precise timer interrupts (less than 1 tick) to benefit the OS scheduler. KVM now exposes this feature to guests.

Persistent Device Naming

This feature stores the mapping of device names (for example, **sda**, **sdb**, and others) and persistent device names (provided by **udev** in **/dev/disk/by-*/**) to kernel messages. This allows users to identify a device from kernel messages. The kernel **/dev/kmsg** log, which can be displayed with the **dmesg** command, now shows the messages for the symbolic links, which **udev** has created for kernel devices. These messages are displayed in the following format:

```
udev-alias: <device_name> (<symbolic_link> <symbolic link> …)
```

Any log analyzer can display these messages, which are also saved in **/var/log/messages** via **syslog**.

New linuxptp Package

The linuxptp package, included in Red Hat Enterprise Linux 6.4 as a Technology Preview, is an implementation of the Precision Time Protocol (PTP) according to IEEE standard 1588 for Linux. The dual design goals are to provide a robust implementation of the standard and to use the most relevant and modern Application Programming Interfaces (API) offered by the Linux kernel. Supporting legacy APIs and other platforms is not a goal.

Transparent Hugepages Documentation

Documentation for transparent hugepages has been added to the following file:

```
/usr/share/doc/kernel-doc-<version>/Documentation/vm/transhuge.txt
```

State of Support for Dump Targets

In Red Hat Enterprise Linux 6.4, the **/usr/share/doc/kexec-tools-2.0.0/kexec-kdump-howto.txt** file provides a comprehensive list of supported, unsupported, and unknown dump targets under section "Dump Target support status".

CHAPTER 3. DEVICE DRIVERS

The *Device Drivers* chapter has been moved to the Red Hat Enterprise Linux Technical Notes, located at:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/6.4_Technical_Notes/ch-device_drivers.html

CHAPTER 4. NETWORKING

HAProxy

HAProxy is a stand-alone, Layer 7, high-performance network load balancer for TCP and HTTP-based applications which can perform various types of scheduling based on the content of the HTTP requests. Red Hat Enterprise Linux 6.4 introduces the haproxy package as a Technology Preview.

Mellanox SR-IOV Support

Single Root I/O Virtualization (SR-IOV) is now supported as a Technology Preview in the Mellanox **libmlx4** library and the following drivers:

- **mlx_core**

- **mlx4_ib** (InfiniBand protocol)

- **mlx_en** (Ethernet protocol)

CHAPTER 5. AUTHENTICATION AND INTEROPERABILITY

SSSD Fully Supported Features
A number of features introduced in Red Hat Enterprise Linux 6.3 are now fully supported in Red Hat Enterprise Linux 6.4. Specifically:

- support for central management of SSH keys,

- SELinux user mapping,

- and support for automount map caching.

New SSSD Cache Storage Type
Kerberos version 1.10 added a new cache storage type, **DIR:**, which allows Kerberos to maintain Ticket Granting Tickets (TGTs) for multiple Key Distribution Centers (KDCs) simultaneously and auto-select between them when negotiating with Kerberos-aware resources. In Red Hat Enterprise Linux 6.4, SSSD has been enhanced to allow you to select the **DIR:** cache for users that are logging in via SSSD. This feature is introduced as a Technology Preview.

Adding AD-based Trusted Domains to `external` Groups
In Red Hat Enterprise Linux 6.4, the **ipa group-add-member** command allows you to add members of Active Directory-based trusted domains to groups marked as **external** in Identity Management. These members may be specified by their name using domain- or UPN-based syntax, for example **AD\UserName** or **AD\GroupName**, or **User@AD.Domain**. When specified in this form, members are resolved against Active Directory-based trusted domain's Global Catalog to obtain their Security Identifier (SID) value.

Alternatively, an SID value could be specified directly. In this case, the **ipa group-add-member** command will only verify that the domain part of the SID value is one of the trusted Active Directory domains. No attempt will be done to verify validity of the SID within the domain.

It is recommended to use user or group name syntax to specify external members rather than providing their SID values directly.

Auto-renew Identity Management Subsystem Certificates
The default validity period for a new Certificate Authority is 10 years. The CA issues a number of certificates for its subsystems (OCSP, audit log, and others). Subsystem certificates are normally valid for 2 years. If the certificates expire, the CA does not start up or does not function properly. Therefore, in Red Hat Enterprise Linux 6.4, Identity Management servers are capable of automatically renewing their subsystem certificates. The subsystem certificates are tracked by **certmonger**, which automatically attempts to renew the certificates before they expire.

Automatic Configuration of OpenLDAP Client Tools on Clients Enrolled in Identity Management
In Red Hat Enterprise Linux 6.4, OpenLDAP is automatically configured with the default LDAP URI, a Base DN, and a TLS certificate during Identity Management client installation. This improves user experience when performing LDAP searches to Identity Management Directory Server.

PKCS#12 Support for `python-nss`
The python-nss package, which provides Python bindings for Network Security Services (NSS) and the Netscape Portable Runtime (NSPR), has been updated to add PKCS #12 support.

Full Persistent Search for DNS
LDAP in Red Hat Enterprise Linux 6.4 includes support for persistent search for both zones and their resource records. Persistent search allows the **bind-dyndb-ldap** plug-in to be immediately informed

about all changes in an LDAP database. It also decreases network bandwidth usage required by repeated polling.

New CLEANALLRUV Operation

Obsolete elements in the Database Replica Update Vector (RUV) can be removed with the **CLEANRUV** operation, which removes them on a single supplier or master. Red Hat Enterprise Linux 6.4 adds a new **CLEANALLRUV** operation which can remove obsolete RUV data from all replicas and needs to be run on a single supplier/master only.

samba4 Libraries Updated

The **samba4** libraries (provided by the samba4-libs package) have been upgraded to the latest upstream version to improve interoperability with Active Directory (AD) domains. SSSD now uses the **libndr-krb5pac** library to parse the Privilege Attribute Certificate (PAC) issued by an AD Key Distribution Center (KDC). Additionally, various improvements have been made to the Local Security Authority (LSA) and Net Logon services to allow verification of trust from a Windows system. For information on the introduction of Cross Realm Kerberos Trust functionality, which depends on samba4 packages, refer to the section called "Cross Realm Kerberos Trust Functionality in Identity Management".

> **WARNING**
>
> If you upgrade from Red Hat Enterprise Linux 6.3 to Red Hat Enterprise Linux 6.4 and you have Samba in use, make sure to uninstall the samba4 package to avoid conflicts during the upgrade.

Because the Cross Realm Kerberos Trust functionality is considered a Technology Preview, selected **samba4** components are considered to be a Technology Preview. For more information on which Samba packages are considered a Technology Preview, refer to Table 5.1, "Samba4 Package Support".

Table 5.1. Samba4 Package Support

Package Name	New Package in 6.4?	Support Status
samba4-libs	No	Technology Preview, except functionality required by OpenChange
samba4-pidl	No	Technology Preview, except functionality required by OpenChange
samba4	No	Technology Preview
samba4-client	Yes	Technology Preview
samba4-common	Yes	Technology Preview
samba4-python	Yes	Technology Preview

Package Name	New Package in 6.4?	Support Status
samba4-winbind	Yes	Technology Preview
samba4-dc	Yes	Technology Preview
samba4-dc-libs	Yes	Technology Preview
samba4-swat	Yes	Technology Preview
samba4-test	Yes	Technology Preview
samba4-winbind-clients	Yes	Technology Preview
samba4-winbind-krb5-locator	Yes	Technology Preview

Cross Realm Kerberos Trust Functionality in Identity Management

The Cross Realm Kerberos Trust functionality provided by Identity Management is included as a Technology Preview. This feature allows to create a trust relationship between an Identity Management and an Active Directory domain. This means that users from the AD domain can access resources and services from the Identity Management domain with their AD credentials. No data needs to be synchronized between the Identity Management and AD domain controllers; AD user are always authenticated against the AD domain controller and information about users is looked up without the need for synchronization.

This feature is provided by the optional ipa-server-trust-ad package. This package depends on features which are only available in **samba4**. Because samba4-* packages conflicts with the corresponding samba-* packages, all samba-* packages must be removed before ipa-server-trust-ad can be installed.

When the ipa-server-trust-ad package is installed, the `ipa-adtrust-install` command must be run on all Identity Management servers and replicas to enable Identity Management to handle trusts. When this is done a trust can be established on the command line using the `ipa trust-add` or the WebUI. For more information, refer to section *Integrating with Active Directory Through Cross-Realm Kerberos Trusts* in the *Identity Management Guide* on https://access.redhat.com/site/documentation/Red_Hat_Enterprise_Linux/.

Posix Schema Support for 389 Directory Server

Windows Active Directory (AD) supports the POSIX schema (RFC 2307 and 2307bis) for user and group entries. In many cases, AD is used as the authoritative source of user and group data, including POSIX attributes. With Red Hat Enterprise Linux 6.4, Directory Server Windows Sync no longer ignores these attributes. Users are now able to synchronize POSIX attributes with Windows Sync between AD and 389 Directory Server.

> **NOTE**
>
> When adding new user and group entries to the Directory Server, the POSIX attributes are not synced to AD. Adding new user and group entries to AD will synchronize to the Directory Server, and modifying attributes will synchronize them both ways.

CHAPTER 6. SECURITY

Treating Matches Authoritatively in Look Ups of sudoers Entries

The **sudo** utility is able to consult the **/etc/nsswitch.conf** file for sudoers entries and look them up in files or using LDAP. Previously, when a match was found in the first database of sudoers entries, the look up operation still continued in other databases (including files). In Red Hat Enterprise Linux 6.4, an option was added to the **/etc/nsswitch.conf** file that allows users to specify a database after which a match of a sudoers entry is sufficient. This eliminates the need to query any other databases; thus, improving the performance of sudoers entry look ups in large environments. This behavior is not enabled by default and must be configured by adding the **[SUCCESS=return]** string after a selected database. When a match is found in a database that directly precedes this string, no other databases are queried.

Additional Password Checks for `pam_cracklib`

The **pam_cracklib** module has been updated to add multiple new password strength checks:

- Certain authentication policies do not allow passwords which contain long continuous sequences such as "abcd" or "98765". This update introduces the possibility to limit the maximum length of these sequences by using the new **maxsequence** option.

- The **pam_cracklib** module now allows to check whether a new password contains the words from the GECOS field from entries in the **/etc/passwd** file. The GECOS field is used to store additional information about the user, such as the user's full name or a phone number, which could be used by an attacker for an attempt to crack the password.

- The **pam_cracklib** module now allows to specify the maximum allowed number of consecutive characters of the same class (lowercase, uppercase, number and special characters) in a password via the **maxrepeatclass** option.

- The **pam_cracklib** module now supports the **enforce_for_root** option, which enforces complexity restrictions on new passwords for the root account.

Size Option for tmpfs Polyinstantiation

On a system with multiple tmpfs mounts, it is necessary to limit their size to prevent them from occupying all of the system memory. PAM has been updated to allow users to specify the maximum size of the tmpfs file system mount when using tmpfs polyinstantiation by using the **mntopts=size=<size>** option in the **/etc/namespace.conf** configuration file.

Locking Inactive Accounts

Certain authentication policies require support for locking of an account that is not used for certain period of time. Red Hat Enterprise Linux 6.4 introduces an additional function to the **pam_lastlog** module, which allows users to lock accounts after a configurable number of days.

New Modes of Operation for `libica`

The **libica** library, which contains a set of functions and utilities for accessing the IBM eServer Cryptographic Accelerator (ICA) hardware on IBM System z, has been modified to allow usage of new algorithms that support the Message Security Assist Extension 4 instructions in the Central Processor Assist for Cryptographic Function (CPACF). For the DES and 3DES block ciphers, the following modes of operation are now supported:

- Cipher Block Chaining with Ciphertext Stealing (CBC-CS)

- Cipher-based Message Authentication Code (CMAC)

For the AES block cipher, the following modes of operation are now supported:

- Cipher Block Chaining with Ciphertext Stealing (CBC-CS)

- Counter with Cipher Block Chaining Message Authentication Code (CCM)

- Galois/Counter (GCM)

This acceleration of complex cryptographic algorithms significantly improves the performance of IBM System z machines.

Optimization of, and Support for, the `zlib` Compression Library for System z

The zlib library, a general-purpose lossless data compression library, has been updated to improve compression performance on IBM System z.

Fallback Firewall Configuration

The **iptables** and **ip6tables** services now provide the ability to assign a fallback firewall configuration if the default configurations cannot be applied. If applying of the firewall rules from **/etc/sysconfig/iptables** fails, the fallback file is applied if it exists. The fallback file is named **/etc/sysconfig/iptables.fallback** and uses the **iptables-save** file format (same as **/etc/sysconfig/iptables**). If application of the fallback file also fails, there is no further fallback. To create a fallback file, use the standard firewall configuration tools and rename or copy the file to the fallback file. Use the same process for the **ip6tables** service, only replace all occurrences of "iptables" with "ip6tables".

CHAPTER 7. SUBSCRIPTION MANAGEMENT

7.1. SUBSCRIPTION MANAGER

String Updates

In Red Hat Enterprise Linux 6.4, several strings have been renamed in Subscription Manager:

- *subscribe* was renamed to *attach*

- *auto-subscribe* was renamed to *auto-attach*

- *unsubscribe* was renamed to *remove*

- *consumer* was renamed to *system* or *unit*

Testing Proxy Connection

The Proxy Configuration dialog now allows users to test a connection to a proxy after entering a value.

Subscribe or Unsubscribe Multiple Entitlements

Subscription Manager is now able to subscribe (attach) or unsubscribe (remove) multiple entitlements using their serial numbers at once.

Activation Keys Support in the GUI

The Subscription Manager graphical user interface now allows you to register a system using an *activation key*. Activation keys allow users to preconfigure subscriptions for a system before it is registered.

Registering Against External Servers

Support for the selection of a remote server during the registration of a system is now supported in Subscription Manager. The Subscription Manager user interface provides an option to choose a URL of a server to register against, together with a port and a prefix, during the registration process. Additionally, when registering on the command line, the `--serverurl` option can be used to specify the server to register against. For more information about this feature, refer to the section *Registering, Unregistering, and Reregistering a System* in the *Subscription Management Guide*.

Usability Changes in the GUI

The Subscription Manager GUI has been enhanced with various changes based on customer feedback.

7.2. SUBSCRIPTION ASSET MANAGER

Installation on Offline Systems

Subscription Asset Manager is now available as an ISO image and can be obtained from Content Delivery Network and Red Hat Network. It is therefore possible to install Subscription Asset Manager on offline systems.

Reduced System Registration Workload

It is now possible to configure a kickstart file with instructions to connect to Subscription Asset Manager and to automatically register and subscribe the system. This significantly reduces workloads of registering a large number of systems.

Migration

Red Hat Enterprise Linux 6.4 provides subscription-manager which includes the **rhn-migrate-classic-to-rhsm** script. The script has the `--serverurl` parameter that allows the user to point the system to an existing or on-premise installation of Subscription Asset Manager, and automatically migrates the system

to use Subscription Asset Manager for its content.

> **NOTE**
>
> For more information about the 1.2 release of Subscription Asset Manager, refer to the Red Hat Subscription Asset Manager 1.2 Release Notes located at:
>
> https://access.redhat.com/site/documentation/en-US/Red_Hat_Subscription_Asset_Manager/1.2/html-single/Release_Notes/index.html

CHAPTER 8. VIRTUALIZATION

8.1. KVM

virtio-SCSI

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 approximately 25 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 ~25-device limit of virtio-blk.

virtio-scsi was introduced in Red Hat Enterprise Linux 6.3 as Technology Preview and is being promoted to fully supported in Red Hat Enterprise Linux 6.4. Windows guests (excluding Windows XP) are also supported with the latest virtio-win drivers.

Support for Intel's Next-generation Core Processor

Red Hat Enterprise Linux 6.4 adds support for Intel's next-generation Core processor to **qemu-kvm** so that KVM guests can utilize new features this processor provides, most important of which are: Advanced Vector Extensions 2 (AVX2), Bit-Manipulation Instructions 1 (BMI1), Bit-Manipulation Instructions 2 (BMI2), Hardware Lock Elision (HLE), Restricted Transactional Memory (RTM), Process-Context Identifier (PCID), Invalidate Process-Context Identifier (INVPCID), Fused Multiply-Add (FMA), Big-Endian Move instruction (MOVBE), F Segment and G Segment BASE instruction (FSGSBASE), Supervisor Mode Execution Prevention (SMEP), Enhanced REP MOVSB/STOSB (ERMS).

Support for AMD Opteron 4xxx Series CPU

The AMD Opteron 4xxx series processor is now supported by **qemu-kvm**. This allows new features of this processor series to be exposed to KVM guests, such as: the F16C instruction set, Trailing Bit Manipulation, Bit-Manipulation Instructions 1 (BMI1) decimate functions, and the Fused Multiply-Add (FMA) instruction set.

Guest Live Migration Using USB Forwarding via SPICE

In Red Hat Enterprise Linux 6.4, KVM supports live migration of guests using USB forwarding via SPICE, while maintaining existing USB device redirection for all configured devices.

Live Migration of Guests Using USB Devices

In Red Hat Enterprise Linux 6.4, KVM supports live migration of guests with USB devices. The following devices are supported: Enhanced Host Controller Interface (EHCI) and Universal Host Controller Interface (UHCI) local passthrough and emulated devices such as storage devices, mice, keyboards, hubs, and others.

QEMU Guest Agent Updated

The QEMU guest agent (provided by the qemu-guest-agent package) is now fully supported in Red Hat Enterprise Linux 6.4. It has been updated to upstream version 1.1, and includes the following notable enhancements and bug fixes:

- The **guest-suspend-disk** and **guest-suspend-ram** commands can now be used to suspend to RAM or to disk on a Windows system.

- The **guest-network-get-interfaces** command can now be used to acquire network interface information in Linux.

- This update provides file system freeze support improvements and fixes.

- This update includes various documentation fixes and small improvements.

Paravirtualized End-of-Interrupt Indication (PV-EOI)

Hosts and guests running Red Hat Enterprise Linux 6.3 and older require two VM exits (context switches from a VM to a Hypervisor) for each interrupt: one to inject the interrupt, and another to signal the end of the interrupt. When both host and guest systems are updated to Red Hat Enterprise Linux 6.4 or newer, they can negotiate a paravirtualized end-of-interrupt feature and only require one switch per interrupt. Consequently, using Red Hat Enterprise Linux 6.4 or newer as both a host and a guest, number of exits is reduced by half for interrupt-intensive workloads, such as incoming network traffic with a virtio network device. This leads to significant reduction in host CPU utilization for such workloads. Note that only edge interrupts are enhanced: for example e1000 networking uses level interrupts and was not improved.

Configurable Sound Pass-through

A sound device can now be detected as a **microphone** or a **speaker** in the guest system (in addition to being detected as **line-in** and **line-out**). Sound devices can now function properly in guest applications that accept only certain types of input for voice recording and audio.

8.2. HYPER-V

Inclusion of, and Guest Installation Support for, Microsoft Hyper-V Drivers

Integrated Red Hat Enterprise Linux guest installation, and Hyper-V para-virtualized device support in Red Hat Enterprise Linux 6.4 on Microsoft Hyper-V allows users to run Red Hat Enterprise Linux 6.4 as a guest on top of Microsoft Hyper-V hypervisors. The following Hyper-V drivers and a clock source have been added to the kernel shipped with Red Hat Enterprise Linux 6.4:

- a network driver (**hv_netvsc**)

- a storage driver (**hv_storvsc**)

- an HID-compliant mouse driver (**hid_hyperv**)

- a VMbus driver (**hv_vmbus**)

- a util driver (**hv_util**)

- an IDE disk driver (**ata_piix**)

- a balloon driver (**hv_balloon**)

- a clock source (i386, AMD64/Intel 64: **hyperv_clocksource**)

Red Hat Enterprise Linux 6.4 also includes support for Hyper-V as a clock source and a guest Hyper-V Key-Value Pair (KVP) daemon (**hypervkvpd**) that passes basic information, such as the guest IP, the FQDN, OS name, and OS release number, to the host through VMbus. An IP injection functionality is also provided which allows you to change the IP address of a guest from the host via the **hypervkvpd** daemon.

Hyper-V `balloon` Driver

On Red Hat Enterprise Linux 6.4 guests, the **balloon** driver, a basic driver for the dynamic memory management functionality supported on Hyper-V hosts, was added. The **balloon** driver is used to dynamically remove memory from a virtual machine. In the current implementation of the balloon driver for Linux, only the ballooning functionality is implemented, not the hot-add functionality.

8.3. VMWARE ESX

VMware PV Drivers

The VMware para-virtualized drivers have been updated to provide a seamless out-of-the-box experience when running Red Hat Enterprise Linux 6.4 in VMware ESX. The Anaconda installer has also been updated to list the drivers during the installation process. The following drivers have been updated:

- a network driver (**vmxnet3**)

- a storage driver (**vmw_pvscsi**)

- a memory ballooning driver (**vmware_balloon**)

- a mouse driver (**vmmouse_drv**)

- a video driver (**vmware_drv**)

CHAPTER 9. CLUSTERING

Support for IBM iPDU Fence Device
Red Hat Enterprise Linux 6.4 adds support for the IBM iPDU fence device. For more information on the parameters of this fence device, refer to the *Fence Device Parameters* appendix in the Red Hat Enterprise Linux 6 *Cluster Administration* guide.

Support for Eaton Network Power Controller Fence Device
Red Hat Enterprise Linux 6.4 adds support for **fence_eaton_snmp**, the fence agent for the Eaton over SNMP network power switch. For more information on the parameters of this fence agent, refer to the *Fence Device Parameters* appendix in the Red Hat Enterprise Linux 6 *Cluster Administration* guide.

New keepalived Package
Red Hat Enterprise Linux 6.4 includes the keepalived package as a Technology Preview. The keepalived package provides simple and robust facilities for load-balancing and high-availability. The load-balancing framework relies on the well-known and widely used Linux Virtual Server kernel module providing Layer 4 network load-balancing. The **keepalived** daemon implements a set of health checkers for load-balanced server pools according to their state. The keepalived daemon also implements the Virtual Router Redundancy Protocol (VRRP), allowing router or director failover to achieve high availability.

Watchdog Recovery
New **fence_sanlock** and **checkquorum.wdmd** fence agents, included in Red Hat Enterprise Linux 6.4 as a Technology Preview, provide new mechanisms to trigger the recovery of a node via a watchdog device. Tutorials on how to enable this Technology Preview will be available at https://fedorahosted.org/cluster/wiki/HomePage.

Support for VMDK-based Storage
Red Hat Enterprise Linux 6.4 adds support for clusters utilizing VMware's VMDK (Virtual Machine Disk) disk image technology with the multi-writer option. This allows you, for example, to use VMDK-based storage with the multi-writer option for clustered file systems such as GFS2.

CHAPTER 10. STORAGE

Support of 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 can improve the scalability and performance of NFS servers for several common workloads.

pNFS defines 3 different storage protocols or layouts: files, objects and blocks. The Red Hat Enterprise Linux 6.4 NFS client supports the files layout protocol.

To enable the pNFS file-layout client in Red Hat Enterprise Linux, use the **-o v4.1** option when mounting a file system on a pNFS-capable server.

When the server is pNFS-enabled, the **nfs_layout_nfsv41_files** kernel module is automatically loaded on the first mount. Use the following command to verify that this module was loaded:

```
~]$ lsmod | grep nfs_layout_nfsv41_files
```

For more information on pNFS, refer to http://www.pnfs.com/.

XFS Online Discard Support
An online discard operation performed on a mounted file system discards blocks which are not in use by the file system. Online discard operations are now supported on XFS file systems. For more information, refer to the section *Discard Unused Blocks* in the Red Hat Enterprise Linux 6 *Storage Administration Guide*.

LVM Support for Micron PCIe SSD
In Red Hat Enterprise Linux 6.4, LVM adds support for Micron PCIe Solid State Drives (SSDs) as devices that may form a part of a Volume Group.

LVM Support for 2-way Mirror RAID10
LVM is now capable of creating, removing, and resizing RAID10 logical volumes. To create a RAID10 logical volume, like the other RAID types, specify the segment type as follows:

```
~]# lvcreate --type raid10 -m 1 -i 2 -L 1G -n lv vg
```

Note that the **-m** and **-i** arguments behave in the same way they would for other segment types. That is, **-i** is the total number of stripes while **-m** is the number of (additional) copies (that is, **-m 1 -i 2** gives 2 stripes on top of 2-way mirrors).

Set Up and Manage SCSI Persistent Reservations Through Device Mapper Devices
Previously, to set up persistent reservations on multipath devices, it was necessary to set it up on all of the path devices. If a path device was later added, it was necessary to manually add reservations to that path. Red Hat Enterprise Linux 6.4 adds the ability to set up and manage SCSI persistent reservations through device mapper devices with the **mpathpersist** command. When path devices are added, persistent reservations are set up on those devices as well.

CHAPTER 11. COMPILER AND TOOLS

SystemTap Updated to Version 1.8

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.

The systemtap package in Red Hat Enterprise Linux 6.4 has been upgraded to upstream version 1.8, which provides a number of bug fixes and enhancements:

- The **@var** syntax is now an alternative language syntax for accessing DWARF variables in **uprobe** and **kprobe** handlers (process, kernel, module).

- SystemTap now mangles local variables to avoid collisions with C headers included by tapsets.

- The SystemTap compile-server and client now support IPv6 networks, for hosts listed in DNS or mDNS.

- The SystemTap runtime (**staprun**) now accepts a **-T** timeout option to allow less frequent wake-ups to poll for low-throughput output from scripts.

- The SystemTap script translator driver (**stap**) now provides the following resource limit options:

  ```
  --rlimit-as=NUM
  --rlimit-cpu=NUM
  --rlimit-nproc=NUM
  --rlimit-stack=NUM
  --rlimit-fsize=NUM
  ```

- SystemTap modules are now smaller and compile faster. The modules' debuginfo is now suppressed by default.

- Bug CVE-2012-0875 (kernel panic when processing malformed DWARF unwind data) is now fixed.

The lscpu and chcpu Utilities

The **lscpu** utility, which displays detailed information about the available CPUs has been updated to include numerous new features. Also, a new utility, **chcpu**, has been added, which allows you to change the CPU state (online/offline, standby/active, and other states), disable and enable CPUs, and configure specified CPUs.

For more information about these utilities, refer to the **lscpu(1)** and **chcpu(8)** man pages.

CHAPTER 12. GENERAL UPDATES

Updated samba Packages

Red Hat Enterprise Linux 6.4 includes rebased samba packages that introduce several bug fixes and enhancements, the most important of which is added support for the SMB2 protocol. SMB2 support can be enabled with the following parameter in the **[global]** section of the **/etc/samba/smb.conf** file:

```
max protocol = SMB2
```

Additionally, Samba now has support for AES Kerberos encryption. AES support has been available in Microsoft Windows operating systems since Windows Vista and Windows Server 2008. It is reported to be the new default Kerberos encryption type since Windows 7. Samba now adds AES Kerberos keys to the keytab it controls. This means that other kerberized services that use the samba keytab and run on the same machine can benefit from AES encryption. In order to use AES session keys (and not only use AES encrypted ticket granting tickets), the samba machine account in Active Directory's LDAP server needs to be manually modified. For more information, refer to the Microsoft Open Specifications Support Team Blog.

With Samba 3.6, it is suggested that if you use the *security = share* mode you should migrate to use *security = user* for a standalone file server or Domain Controller (DC). The *security = share* mode will not be supported in future releases. Refer to the **smb.conf(5)** manpage for more details on *security = user* and read the ACL documentation for permission control on files and directories.

> **WARNING**
>
> The updated samba packages also change the way ID mapping is configured. Users are advised to modify their existing Samba configuration files.
>
> Note that several Trivial Database (TDB) files have been updated and the printing support has been rewritten to use the actual registry implementation. This means that all TDB files are upgraded as soon as you start the new version of **smbd**. You cannot downgrade to an older Samba 3.x version unless you have backups of the TDB files.
>
> For more information about these changes, refer to the Release Notes for Samba 3.6.0.

New SciPy Package

Red Hat Enterprise Linux 6.4 includes a new scipy package. The SciPy package provides software for mathematics, science, and engineering. The NumPy package, which is designed to manipulate large multi-dimensional arrays of arbitrary records, is the core library for SciPy. The SciPy library is built to work with NumPy arrays and provides various efficient numerical routines, for example routines for numerical integration and optimization.

TLS v1.1 Support in NSS

The nss and nss-util packages have been upgraded to upstream version 3.14 to provide, among other features, support for TLS version 1.1. As well, the nspr package has been rebased to version 4.9.2. For more information, refer to the NSS 3.14 Release Notes.

Embedded Valgrind `gdbserver`

The valgrind package has been upgraded to upstream version 3.8.1. This updated version contains, among other enhancements and bug fixes, an embedded **gdbserver**. For more information, refer to the *Valgrind* chapter and the *Changes in Valgrind 3.8.1* appendix in the *Red Hat Developer Toolset 1.1 User Guide*.

New libjpeg-turbo Packages

Red Hat Enterprise Linux 6.4 includes a new set of packages: libjpeg-turbo. These packages replace the traditional libjpeg packages, and provide the same functionality and API as libjpeg but better performance.

New redhat-lsb-core Package

When installing the redhat-lsb package, a large number of dependencies are pulled into the system to meet the LSB standard. Red Hat Enterprise Linux 6.4 provides a new redhat-lsb-core subpackage which allows you to easily fetch only the minimal set of required packages by installing the redhat-lsb-core package.

createrepo Utility Updated

The **createrepo** utility has been updated to the latest upstream version, which significantly reduces memory usage and adds multitasking support via the **--workers** option.

APPENDIX A. COMPONENT VERSIONS

This appendix is a list of components and their versions in the Red Hat Enterprise Linux 6.4 release.

Table A.1. Component Versions

Component	Version
Kernel	2.6.32-358
QLogic `qla2xxx` driver	8.04.00.08.06.4-k
QLogic ql2xxx firmware	ql23xx-firmware-3.03.27-3.1
	ql2100-firmware-1.19.38-3.1
	ql2200-firmware-2.02.08-3.1
	ql2400-firmware-5.08.00-1
	ql2500-firmware-5.08.00-1
Emulex `lpfc` driver	8.3.5.86.1p
iSCSI initiator utils	iscsi-initiator-utils-6.2.0.873-2
DM-Multipath	device-mapper-multipath-0.4.9-64
LVM	lvm2-2.02.98-9

APPENDIX B. REVISION HISTORY

Revision 1.3-7   **Fri Mar 7 2014**   **Eliška Slobodová**
  Updated a note about pNFS.

Revision 1.3-4   **Tue Feb 18 2014**   **Eliška Slobodová**
  Added a note about Mellanox SR-IOV support.

Revision 1.3-3   **Wed Jan 15 2014**   **Eliška Slobodová**
  Updated a note about the Hyper-V balloon driver.

Revision 1.3-2   **Mon Feb 25 2013**   **Martin Prpič**
  Added Subscription Asset Manager release notes.

Revision 1.2-1   **Thu Feb 21 2013**   **Martin Prpič**
  Release of the Red Hat Enterprise Linux 6.4 Release Notes.

Revision 1.1-14   **Wed Dec 4 2012**   **Martin Prpič**
  Release of the Red Hat Enterprise Linux 6.4 Beta Release Notes.