



Red Hat Enterprise Linux 7 7.0 Release Notes

Release Notes for Red Hat Enterprise Linux 7

Red Hat Engineering Content Services

Red Hat Enterprise Linux 7 7.0 Release Notes

Release Notes for Red Hat Enterprise Linux 7

Red Hat Engineering Content Services

Legal Notice

Copyright © 2014 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, MetaMatrix, Fedora, the Infinity Logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

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

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

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

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

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 documents the major features and enhancements implemented in Red Hat Enterprise Linux 7 and the known issues in this 7.0 release. For detailed information regarding the changes between Red Hat Enterprise Linux 6 and 7, consult the Migration Planning Guide. Acknowledgements Red Hat Global Support Services would like to recognize Sterling Alexander and Michael Everette for their outstanding contributions in testing Red Hat Enterprise Linux 7.

Table of Contents

| | |
|---|-----------|
| Introduction | 4 |
| Part I. New Features | 5 |
| Chapter 1. Architectures | 6 |
| Chapter 2. Installation and Booting | 7 |
| 2.1. Installer | 7 |
| 2.2. Boot Loader | 10 |
| Chapter 3. Storage | 11 |
| LIO kernel Target Subsystem | 11 |
| LVM Cache | 11 |
| Storage Array Management with libStorageMgmt API | 11 |
| Support for LSI Syncro | 11 |
| LVM Application Programming Interface | 11 |
| DIF/DIX Support | 11 |
| Support of Parallel NFS | 12 |
| Chapter 4. File Systems | 13 |
| Support of XFS File System | 13 |
| Support of Btrfs File System | 13 |
| Fast Block Devices Caching Slower Block Devices | 13 |
| Chapter 5. Kernel | 14 |
| Dynamic kernel Patching | 14 |
| Support for Large crashkernel Sizes | 14 |
| Crashkernel With More Than 1 CPU | 14 |
| Swap Memory Compression | 14 |
| NUMA-Aware Scheduling and Memory Allocation | 14 |
| APIC Virtualization | 14 |
| vmcp Built in the Kernel | 14 |
| Hardware Error Reporting Mechanism | 14 |
| Full DynTick Support | 15 |
| Blacklisting kernel Modules | 15 |
| dm-erra Target | 15 |
| libhugetlbfs Support for IBM System z | 15 |
| AMD Microcode and AMD Opteron Support | 15 |
| Available Memory for /proc/meminfo | 16 |
| Chapter 6. Virtualization | 17 |
| 6.1. Kernel-Based Virtualization | 17 |
| 6.2. Xen | 21 |
| 6.3. Hyper-V | 21 |
| 6.4. VMware | 21 |
| Chapter 7. Linux Containers with Docker Format | 22 |
| 7.1. Components of Docker Containers | 22 |
| 7.2. Advantages of Using Docker | 23 |
| 7.3. Comparison with Virtual Machines | 23 |
| 7.4. Using Docker on Red Hat Enterprise Linux 7 | 24 |
| Chapter 8. System and Services | 25 |
| systemd | 25 |
| Chapter 9. Clustering | 26 |
| 9.1. Pacemaker Cluster Manager | 26 |

| | |
|---|-----------|
| 9.1. Pacemaker Cluster Manager | 20 |
| 9.2. Keepalived and HAProxy Replace Piranha as Load Balancer | 26 |
| 9.3. High Availability Administration | 27 |
| 9.4. New Resource Agents | 27 |
| Chapter 10. Compiler and Tools | 28 |
| 10.1. GCC Toolchain | 28 |
| 10.2. GLIBC | 28 |
| 10.3. GDB | 29 |
| 10.4. Performance Tools | 30 |
| 10.5. Programming Languages | 33 |
| Chapter 11. Networking | 35 |
| NetworkManager | 35 |
| Networking Team Driver | 35 |
| Precision Time Protocol | 35 |
| chrony Suite | 35 |
| Dynamic Firewall Daemon, firewalld Suite | 35 |
| DNSSEC | 35 |
| DDoS Protection | 36 |
| Support for 40 Gigabit NICs | 36 |
| WiGig 60 GHz Band (IEEE 802.11ad) | 36 |
| Network Namespaces | 36 |
| Trusted Network Connect | 36 |
| SR-IOV Functionality in the qlcnict Driver | 36 |
| FreeRADIUS 3.0.1 | 36 |
| OpenLMI | 37 |
| Chapter 12. Resource Management | 38 |
| Control Groups | 38 |
| Chapter 13. Authentication and Interoperability | 39 |
| Improved Identity Management Cross-Realm Trusts to Active Directory | 39 |
| Support of POSIX User and Group IDs In Active Directory | 39 |
| Use of AD and LDAP sudo Providers | 39 |
| Support of CA-Less Installations | 39 |
| FreeIPA GUI Improvements | 39 |
| Reclaiming IDs of Deleted Replicas | 39 |
| Re-Enrolling Clients Using Existing Keytab Files | 40 |
| Prompt for DNS | 40 |
| Enhanced SSHFP DNS Records | 40 |
| Filtering Groups by Type | 40 |
| Improved Integration with the External Provisioning Systems | 40 |
| CRL and OCSP DNS Name in Certificate Profiles | 40 |
| Certificates Search | 40 |
| Marking Kerberos Service as Trusted for Delegation of User Keys | 41 |
| Samba 4.1.0 | 41 |
| Chapter 14. Security | 42 |
| OpenSSH chroot Shell Logins | 42 |
| OpenSSH - Multiple Required Authentications | 42 |
| GSS Proxy | 42 |
| Changes in NSS | 42 |
| New Boolean Names | 42 |
| SCAP Workbench | 42 |
| OSCAP Anaconda Add-On | 43 |
| Chapter 15. Subscription Management | 44 |
| Certificate-Based Entitlements | 44 |

| | |
|--|-----------|
| Chapter 16. Desktop | 45 |
| 16.1. GNOME 3 | 45 |
| 16.2. KDE | 46 |
| Chapter 17. Web Servers and Services | 47 |
| Apache HTTP Server 2.4 | 47 |
| MariaDB 5.5 | 47 |
| PostgreSQL 9.2 | 47 |
| Chapter 18. Red Hat Software Collections | 48 |
| Chapter 19. Documentation | 49 |
| 19.1. Release Documentation | 49 |
| 19.2. Installation and Deployment | 49 |
| 19.3. Security | 50 |
| 19.4. Tools and Performance | 50 |
| 19.5. Clustering and High Availability | 51 |
| 19.6. Virtualization | 51 |
| Chapter 20. Internationalization | 52 |
| 20.1. Red Hat Enterprise Linux 7 International Languages | 52 |
| 20.2. General Changes In Internationalization | 53 |
| 20.3. Input Methods | 54 |
| 20.4. Fonts | 54 |
| 20.5. Language-Specific Changes | 54 |
| Chapter 21. Supportability and Maintenance | 56 |
| ABRT 2.1 | 56 |
| Additional Information on ABRT | 56 |
| Part II. Known Issues | 57 |
| Chapter 22. Installation | 58 |
| Chapter 23. Storage | 63 |
| Chapter 24. Kernel | 65 |
| Chapter 25. Virtualization | 69 |
| Chapter 26. Deployment and Tools | 73 |
| Chapter 27. Clustering | 74 |
| Chapter 28. Networking | 75 |
| Chapter 29. Authentication and Interoperability | 77 |
| Chapter 30. Security | 81 |
| Chapter 31. Entitlement | 82 |
| Chapter 32. Desktop | 83 |
| Revision History | 84 |

Introduction

Red Hat is pleased to announce the availability of Red Hat Enterprise Linux 7, the next generation of Red Hat's comprehensive suite of operating systems, designed for mission-critical enterprise computing and certified by top enterprise software and hardware vendors.

In this release, Red Hat brings together improvements across the server, systems, and the overall Red Hat open source experience. Among others, Red Hat Enterprise Linux 7 introduces:

- ▶ a new boot loader and a fully redesigned graphical installer;
- ▶ the kernel patching utility Technology Preview, which allows users to patch the kernel without rebooting;
- ▶ the Docker environment that allows users to deploy any application as a lightweight and portable container;
- ▶ the Hardware Event Report Mechanism (HERM) that refactors the Error Detection and Correction (EDAC) mechanism of dual in-line memory module (DIMM) error reporting;
- ▶ the OpenLMI project providing a common infrastructure for the management of Linux systems;
- ▶ Red Hat Software Collections that provides a set of dynamic programming languages, database servers, and related packages.

For detailed information about new features, see the respective categories in [Part I, “New Features”](#). Known problems are listed in [Part II, “Known Issues”](#).

If you are upgrading from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7, consult the [Migration Planning Guide](#) for assistance with the migration process.

Capabilities and limits of Red Hat Enterprise Linux 7 as compared to the previous versions of the system are available in the following Knowledge Base article:

<https://access.redhat.com/site/articles/rhel-limits>

Part I. New Features

This part describes new features and major enhancements introduced in Red Hat Enterprise Linux 7.

Chapter 1. Architectures

Red Hat Enterprise Linux 7 is available as a single kit on the following architectures ^[1]:

- ▶ 64-bit AMD
- ▶ 64-bit Intel
- ▶ IBM POWER7
- ▶ IBM System z ^[2]

In this release, Red Hat brings together improvements for servers and systems, as well as for the overall Red Hat open source experience.

[1] Note that the Red Hat Enterprise Linux 7 installation is only supported on 64-bit hardware.

Red Hat Enterprise Linux 7 is able to run 32-bit operating systems, including previous versions of Red Hat Enterprise Linux, as virtual machines.

[2] Note that Red Hat Enterprise Linux 7 supports IBM zEnterprise 196 hardware or later; IBM System z10 mainframe systems are no longer supported and will not boot Red Hat Enterprise Linux 7.

Chapter 2. Installation and Booting

2.1. Installer

The Red Hat Enterprise Linux installer (known as **anaconda**) assists in the installation of Red Hat Enterprise Linux 7. This section of the *Release Notes* provides an overview of the new features implemented in the installer for Red Hat Enterprise Linux 7.

The new installer in Red Hat Enterprise Linux 7 features a wide range of bug fixes and enhancements, including: a fully redesigned graphical installer and major updates to the storage configuration tools.

2.1.1. Installation Methods

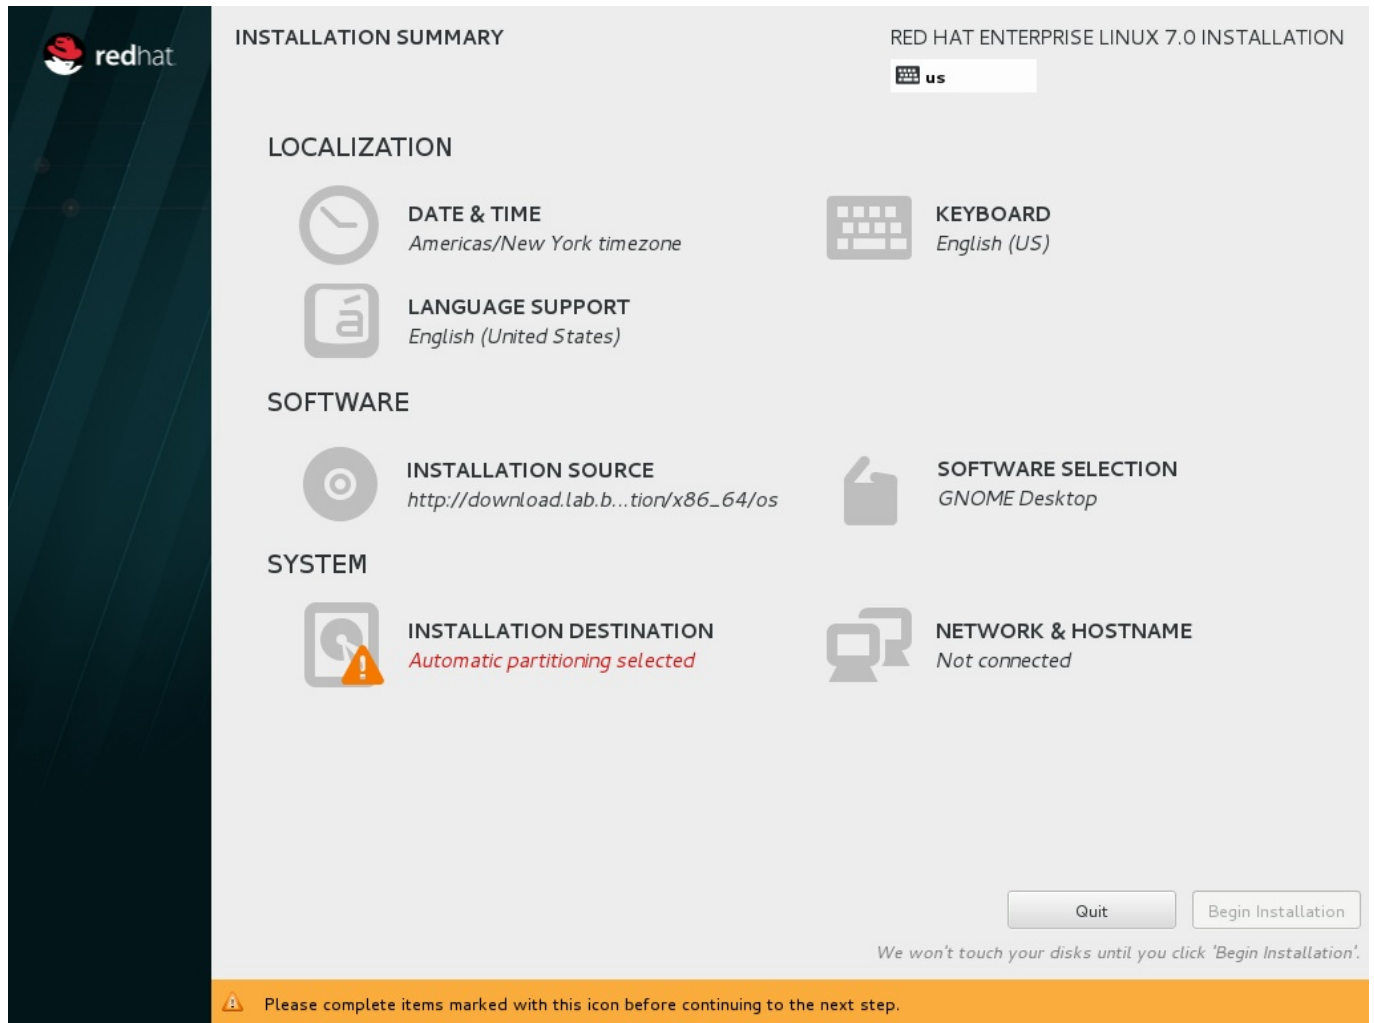
The installer provides three main interfaces to install Red Hat Enterprise Linux:

- the graphic installer,
- the text-based installer,
- and **kickstart**.

2.1.1.1. Graphical Installer

The Red Hat Enterprise Linux graphical installer provides an intuitive graphical user interface to prepare a system for installation. The Red Hat Enterprise Linux 7 graphical installer introduces a brand new user interface designed to make installation quicker and easier.

Previously, the installer was a series of wizard-style screens that required the user to review the settings and click to get to the next screen. The new installer interface provides a central hub that lists groups of configuration options for an installation; the user clicks on the options that need changing, changes them, then initiates the installation.



The new graphical installer also generates automatic default settings where applicable. For example, if the installer detects a network connection, the user's general location is determined with GeoIP and sane suggestions are made for the default keyboard layout, language and timezone.

Additionally, the graphical installer processes some tasks concurrently (for example, storage layout detection) allowing the user to continue configuring the installation using the GUI while the processor-intensive tasks are processed in the background.

2.1.1.2. Text-Based Installer

The text-based installer is provided primarily for systems with limited resources. Red Hat Enterprise Linux 7 features a completely rewritten text-mode installer that provides better support for serial consoles and other limited display interfaces. The text-based installer utilizes the **tmux** utility, making multiple shell terminals available for all installation methods, not just those supporting Linux virtual consoles.

2.1.1.3. Kickstart

Kickstart is an automated installation method that system administrators can use to install Red Hat Enterprise Linux. Using **kickstart**, a single file is created, containing the answers to all the questions that would normally be asked during a typical installation. **Kickstart** in Red Hat Enterprise Linux 7 supports Active Directory host enrollment using the **kickstart** service **realm** **drealmd**.

2.1.2. Plug-In Architecture

The installer in Red Hat Enterprise Linux 7 supports the development of plug-ins that can provide site-specific extensions or customization. Plug-ins can be developed to add additional screens and options to the graphical installer. The plug-in architecture also allows developers to add new **kickstart** commands for system administrators to utilize.

2.1.3. Storage Features and Enhancements

2.1.3.1. Custom Partitioning

The screenshot displays the 'MANUAL PARTITIONING' screen for 'RED HAT ENTERPRISE LINUX 7.0 INSTALLATION'. The left sidebar shows a tree view of the installation structure:

- DATA
- SYSTEM
 - /boot (500 MB, vda1)
 - / (6.86 GB, luks-rhel-root) - **Selected**
 - swap (819 MB, rhel-swap)

At the bottom left, a summary shows: AVAILABLE SPACE: 969.23 kB, TOTAL SPACE: 8.19 GB, and 1 storage device selected.

The right panel shows configuration for the selected 'luks-rhel-root' volume:

- Name: root
- Mount Point: /
- Label: (empty)
- Desired Capacity: 6.866 GB
- Device Type: LVM
- File System: xfs
- Volume Group: rhel (0 B free)

Checkboxes for 'Encrypt' and 'Reformat' are checked. Buttons for 'Update Settings', 'Modify...', and 'Reset All' are present. A note at the bottom states: 'Note: The settings you make on this screen will not be applied until you click on the main menu's 'Begin Installation' button.'

2.1.3.2. Rescanning Storage

The installer does not expose all possible storage tunables in the user interface. In order to accommodate users who require very low-level configuration of their storage, the user can exit the installer to perform their storage configuration. The user can then return to the installer and have it rescan the storage to detect their configuration and present it in the graphical interface.

2.1.3.3. Automatic Partitioning

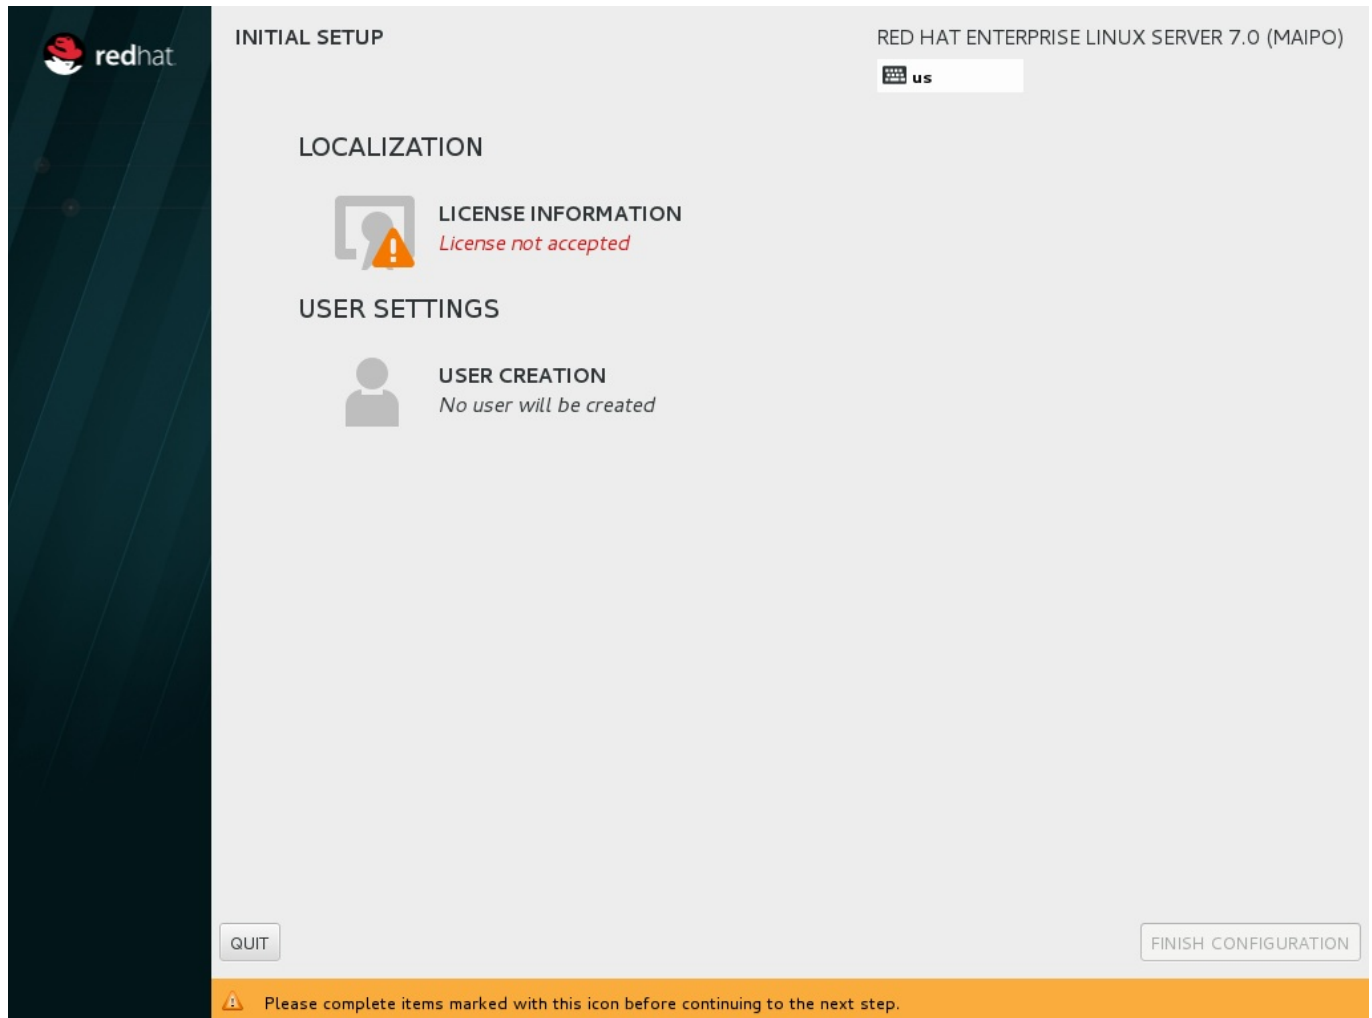
The Red Hat Enterprise Linux 7 installer offers more automatic partitioning choices; for example, LVM, LVM with thin provisioning, BTRFS, or standard partitions.

2.1.3.4. Installation Environment

Previously, the installation environment implemented its own initialization and device discovery tools that were different from the tools used to boot the installed system. In Red Hat Enterprise Linux 7, the installer utilizes the same initialization (**systemd**>) and device discovery tools (**dracut**) as the installed system.

2.1.3.5. Initial System Configuration

After installation, the initial system configuration screens allow for further configuration of Red Hat Enterprise Linux 7 installation. The Initial System Configuration screens in Red Hat Enterprise Linux 7 are also redesigned to match the user experience of the new installer graphical interface. Additionally, some tasks that were traditionally only configurable post-installation (for example, creating the initial user) can now be configured in the installer while the system is being installed.



2.2. Boot Loader

GRUB 2

Red Hat Enterprise Linux 7 includes a new boot loader, GRUB 2, which is more robust, portable, and powerful than its predecessor, GRUB, which is the boot loader that Red Hat Enterprise Linux 6 uses. GRUB 2 provides a number of features and improvements, the most notable of which are:

- In addition to the 64-bit Intel and AMD architectures, GRUB 2 supports a wider variety of platforms, including PowerPC.
- GRUB 2 supports additional firmware types, including BIOS, EFI and OpenFirmware.
- In addition to supporting Master Boot Record (MBR) partition tables, GRUB 2 supports GUID Partition Tables (GPT).
- In addition to the Linux file systems, GRUB 2 also supports non-Linux file systems such as **Apple Hierarchical File System Plus (HFS+)** and Microsoft **NTFS** file system.

Chapter 3. Storage

LIO kernel Target Subsystem

Red Hat Enterprise Linux 7 uses the LIO kernel target subsystem, which is the standard open source SCSI target for block storage, for all of the following storage fabrics: FCoE, iSCSI, iSER (Mellanox InfiniBand), and SRP (Mellanox InfiniBand).

Red Hat Enterprise Linux 6 uses **tgtd**, the SCSI Target Daemon, for iSCSI target support, and only uses LIO, the Linux kernel target, for Fibre-Channel over Ethernet (FCoE) targets via the *fcoe-target-utils* package.

The **targetcli** shell provides the general management platform for the LIO Linux SCSI target.

LVM Cache

Red Hat Enterprise Linux 7 introduces LVM cache as a Technology Preview. This feature allows users to create logical volumes with a small fast device performing as a cache to larger slower devices. Please refer to the **lvmd (8)** manual page for information on creating cache logical volumes.

Note that the following commands are not currently allowed on cache logical volumes:

- ▶ **pvmove**: will skip over any cache logical volume;
- ▶ **lvresize**, **lvreduce**, **lvextend**: cache logical volumes cannot be resized currently;
- ▶ **vgsplit**: splitting a volume group is not allowed when cache logical volumes exist in it.

Storage Array Management with libStorageMgmt API

Red Hat Enterprise Linux 7 introduces storage array management as a Technology Preview. libStorageMgmt is a storage array independent Application Programming Interface (API). It provides a stable and consistent API that allows developers to programmatically manage different storage arrays and utilize the hardware-accelerated features provided. System administrators can also use it as a tool to manually configure storage and to automate storage management tasks with the included Command Line Interface (CLI).

Support for LSI Syncro

Red Hat Enterprise Linux 7 includes code in the **megaraid_sas** driver to enable LSI Syncro CS high-availability direct-attached storage (HA-DAS) adapters. While the **megaraid_sas** driver is fully supported for previously enabled adapters, the use of this driver for Syncro CS is available as a Technology Preview. Support for this adapter will be provided directly by LSI, your system integrator, or system vendor. Users deploying Syncro CS on Red Hat Enterprise Linux 7 are encouraged to provide feedback to Red Hat and LSI. For more information on LSI Syncro CS solutions, please visit <http://www.lsi.com/products/shared-das/pages/default.aspx>.

LVM Application Programming Interface

Red Hat Enterprise Linux 7 features the new LVM application programming interface (API) as a Technology Preview. This API is used to query and control certain aspects of LVM.

Refer to the **lvmd2app.h** header file for more information.

DIF/DIX Support

DIF/DIX is a new addition to the SCSI Standard and a Technology Preview in Red Hat Enterprise Linux 7. DIF/DIX increases the size of the commonly used 512-byte disk block from 512 to 520 bytes, adding the Data Integrity Field (DIF). The DIF stores a checksum value for the data block that is calculated by the Host Bus Adapter (HBA) when a write occurs. The storage device then confirms the checksum on receive, and stores both the data and the checksum. Conversely, when a read occurs, the checksum can be verified by the storage device, and by the receiving HBA.

For more information, refer to the section Block Devices with DIF/DIX Enabled in the [Storage Administration Guide](#).

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 three different storage protocols or layouts: files, objects, and blocks. The Red Hat Enterprise Linux 7 client fully supports the files layout, and the blocks and object layouts are supported as a Technology Preview.

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

Chapter 4. File Systems

Support of XFS File System

The default file system for an **Anaconda**-based installation of Red Hat Enterprise Linux 7 is now **XFS**, which replaces the Fourth Extended Filesystem (**ext4**) used by default in Red Hat Enterprise Linux 6. The **ext4**, **ext3** and **ext2** file systems can be used as alternatives to **XFS**.

XFS is a highly scalable, high-performance file system which was originally designed at Silicon Graphics, Inc. It was created to support file systems up to 16 exabytes (approximately 16 million terabytes), files up to 8 exabytes (approximately 8 million terabytes) and directory structures containing tens of millions of entries. **XFS** supports metadata journaling, which facilitates quicker crash recovery. **XFS** file system can also be defragmented and expanded while mounted and active. Note that it is not possible to shrink XFS file system.

For information about changes between commands used for common tasks in **ext4** and **XFS**, see the Reference Table in the [Installation Guide](#).

Support of Btrfs File System

The **Btrfs** (B-Tree) file system is supported as a Technology Preview in Red Hat Enterprise Linux 7. This file system offers advanced management, reliability, and scalability features. It enables users to create snapshots, it allows for compression and integrated device management.

For more information about the Btrfs Technology Preview, see [Storage Administration Guide](#)

Fast Block Devices Caching Slower Block Devices

LVM provides the ability to have fast block devices act as a cache for slower block devices. This feature is introduced as a Technology Preview in Red Hat Enterprise Linux 7 and allows a PCIe SSD device to act as a cache for direct-attached storage (DAS) or storage area network (SAN) storage, which improves file system performance.

Please refer to the **lvm (8)** manual page for detailed information.

Chapter 5. Kernel

Red Hat Enterprise Linux 7 includes *kernel* version 3.10, which provides a number of new features, the most notable of which are listed below.

Dynamic kernel Patching

Red Hat Enterprise Linux 7 introduces **kpatch**, a dynamic "kernel patching utility", as a Technology Preview. **kpatch** allows users to manage a collection of binary kernel patches which can be used to dynamically patch the kernel without rebooting. Note that **kpatch** is supported to run on AMD64 and Intel 64 architectures only.

Support for Large crashkernel Sizes

Red Hat Enterprise Linux 7 supports the **kdump** crash dumping mechanism on systems with large memory (up to 3TB).

Crashkernel With More Than 1 CPU

Red Hat Enterprise Linux 7 enables booting crashkernel with more than one CPU. This function is supported as a Technology Preview.

Swap Memory Compression

Red Hat Enterprise Linux 7 introduces a new feature, swap memory compression. Swap compression is performed through **zswap**, a thin back end for **frontswap**. Utilizing the swap memory compression technology ensures a significant I/O reduction and performance gains.

NUMA-Aware Scheduling and Memory Allocation

In Red Hat Enterprise Linux 7, the kernel automatically relocates processes and memory between NUMA nodes in the same system, in order to improve performance on systems with non-uniform memory access (NUMA).

APIC Virtualization

Virtualization of Advanced Programmable Interrupt Controller (APIC) registers is supported by utilizing hardware capabilities of new processors to improve virtual machine monitor (VMM) interrupt handling.

vmcp Built in the Kernel

In Red Hat Enterprise Linux 7, the **vmcp** kernel module is built into the kernel. This ensures that the **vmcp** device node is always present, and users can send IBM z/VM hypervisor control program commands without having to load the **vmcp** kernel module first.

Hardware Error Reporting Mechanism

The hardware error reporting mechanisms could previously be problematic because various tools were used to collect errors from different sources with different methods, and different tools were used to report the error events. Red Hat Enterprise Linux 7 introduces Hardware Event Report Mechanism, or HERM. This new infrastructure refactors the Error Detection and Correction (EDAC) mechanism of dual in-line memory module (DIMM) error reporting and also provides new ways to gather system-reported memory errors. The error events are reported to user space in a sequential timeline and single location.

HERM in Red Hat Enterprise Linux 7 also introduces a new user space daemon, **rasdaemon**, which replaces the tools previously included in the *edac-utils* package. The **rasdaemon** catches and handles all Reliability, Availability, and Serviceability (RAS) error events that come from the kernel tracing infrastructure, and logs them. HERM in Red Hat Enterprise Linux 7 also provides the tools to report the errors and is able to detect different types of errors such as burst and sparse errors.

Full DynTick Support

The **nohz_full** boot parameter extends the original tickless kernel feature to an additional case when the tick can be stopped, when the per-cpu **nr_running=1** setting is used. That is, when there is a single runnable task on a CPU's run queue.

Blacklisting kernel Modules

The **modprobe** utility included with Red Hat Enterprise Linux 7 allows users to blacklist kernel modules at installation time. To globally disable autoloading of a module, use this option on the kernel command line:

```
modprobe.blacklist=module
```

For more information on **kpatch**, see <http://rhelblog.redhat.com/2014/02/26/kpatch/>.

dm-era Target

Red Hat Enterprise Linux 7 introduces the dm-era device-mapper target as a Technology Preview. dm-era keeps track of which blocks were written within a user-defined period of time called an "era". Each era target instance maintains the current era as a monotonically increasing 32-bit counter. This target enables backup software to track which blocks have changed since the last backup. It also allows for partial invalidation of the contents of a cache to restore cache coherency after rolling back to a vendor snapshot. The dm-era target is primarily expected to be paired with the dm-cache target.

libhugetlbfs Support for IBM System z

The **libhugetlbfs** library is now supported on IBM System z architecture. The library enables transparent exploitation of large pages in C and C++ programs. Applications and middleware programs can profit from the performance benefits of large pages without changes or recompilations.

AMD Microcode and AMD Opteron Support

AMD provides microcode patch support for processors belonging to AMD processor families 10h, 11h, 12h, 14h, and 15h. Microcode patches contain fixes for processor errata, which ensures that the processor microcode patch level is at the latest level.

One single container file contains all microcode patches for AMD families 10h, 11h, 12h, 14h processors. A separate container file contains patches for AMD family 15h processors.

Note that microcode patches are not incremental, therefore, you only need to make sure you have the latest container file for your AMD processor family. To obtain these microcode patches for your AMD-based platform running Red Hat Enterprise Linux 7:

1. Clone the repository with firmware files.

```
~]$ git clone
git://git.kernel.org/pub/scm/linux/kernel/git/firmware/linux-firmware.git
```

2. Move the AMD microcode files into the **/lib/firmware/** directory. As **root**:

```
~]# cp -r linux-firmware/amd-ucode/ /lib/firmware/
```

Available Memory for `/proc/meminfo`

A new entry to the `/proc/meminfo` file has been introduced to provide the **MemAvailable** field. **MemAvailable** provides an estimate of how much memory is available for starting new applications, without swapping. However, unlike the data provided by the **Cache** or **Free** fields, **MemAvailable** takes into account page cache and also that not all reclaimable memory slabs will be reclaimable due to items being in use.

Chapter 6. Virtualization

6.1. Kernel-Based Virtualization

Improved Block I/O Performance Using `virtio-blk-data-plane`

In Red Hat Enterprise Linux 7, the **`virtio-blk-data-plane`** I/O virtualization functionality is available as a Technology Preview. This functionality extends QEMU to perform disk I/O in a dedicated thread that is optimized for I/O performance.

PCI Bridge

QEMU previously supported only up to 32 PCI slots. Red Hat Enterprise Linux 7 features PCI Bridge as a Technology Preview. This functionality allows users to configure more than 32 PCI devices. Note that hot plugging of devices behind the bridge is not supported.

QEMU Sandboxing

Red Hat Enterprise Linux 7 features enhanced KVM virtualization security through the use of kernel system call filtering, which improves isolation between the host system and the guest.

QEMU Virtual CPU Hot Add Support

QEMU in Red Hat Enterprise Linux 7 features virtual CPU (vCPU) hot add support. Virtual CPUs (vCPUs) can be added to a running virtual machine in order to meet either the workload's demands or to maintain the Service Level Agreement (SLA) associated with the workload. Note that vCPU hot plug is only supported on virtual machines using the **`pc-i440fx-rhel7.0.0`** machine type, the default machine type on Red Hat Enterprise Linux 7.

Multiple Queue NICs

Multiple queue `virtio_net` provides better scalability; each virtual CPU can have a separate transmit or receive queue and separate interrupts that it can use without influencing other virtual CPUs. Note that this feature is only supported on Linux guests.

Multiple Queue `virtio_scsi`

Multiple queue `virtio_scsi` provides better scalability; each virtual CPU can have a separate queue and interrupts that it can use without influencing other virtual CPUs. Note that this feature is only supported on Linux guests.

Page Delta Compression for Live Migration

The KVM live migration feature has been improved by compressing the guest memory pages and reducing the size of the transferred migration data. This feature allows the migration to converge faster.

HyperV Enlightenment in KVM

KVM has been updated with several Microsoft Hyper-V functions; for example, support for Memory Management Unit (MMU) and Virtual Interrupt Controller. Microsoft provides a para-virtualized API between the guest and the host, and by implementing parts of this functionality on the host, and exposing it according to Microsoft specifications, Microsoft Windows guests can improve their performance. Note that these functions are not enabled by default.

EOI Acceleration for High Bandwidth I/O

Red Hat Enterprise Linux 7 utilizes Intel and AMD enhancements to Advanced Programmable Interrupt Controller (APIC) to accelerate end of interrupt (EOI) processing. For older chipsets, Red Hat Enterprise Linux 7 provides para-virtualization options for EOI acceleration.

USB 3.0 Support for KVM Guests

Red Hat Enterprise Linux 7 features improved USB support by adding USB 3.0 host adapter (xHCI) emulation as a Technology Preview.

Microsoft Windows and Windows Server Guest Support

Red Hat Enterprise Linux 7 supports Microsoft Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 guests running inside KVM virtual machines.

I/O Throttling for QEMU Guests

This feature provides I/O throttling, or limits, for QEMU guests' block devices. I/O throttling slows down the processing of disk I/O requests. This slows down one guest disk to reserve I/O bandwidth for other tasks on host. Note that currently it is not possible to throttle virtio-blk-data-plane devices.

Integration of Ballooning and Transparent Huge Pages

Ballooning and transparent huge pages are better integrated in Red Hat Enterprise Linux 7. Balloon pages can be moved and compacted so they can become huge pages.

Pulling System Entropy from Host

A new device, **virtio-rng**, can be configured for guests, which will make entropy available to guests from the host. By default, this information is sourced from the host's **/dev/random** file, but hardware random number generators (RNGs) available on hosts can be used as the source as well.

Bridge Zero Copy Transmit

Bridge zero-copy transmit is a performance feature to improve CPU processing of large messages. The bridge zero-copy transmit feature improves performance from guest to external traffic when using a bridge. Note that this function is disabled by default.

Live Migration Support

Live migration of a guest from a Red Hat Enterprise Linux 6.5 host to a Red Hat Enterprise Linux 7 host is supported.

Discard Support in qemu-kvm

Discard support, using the **fstrim** or **mount -o discard** command, works on a guest after adding **discard='unmap'** to the **<driver>** element in the domain's XML definition. For example:

```
<disk type='file' device='disk'>
  <driver name='qemu' type='raw' discard='unmap' />
  <source file='/var/lib/libvirt/images/vm1.img'>
    ...
  </disk>
```

NVIDIA GPU Device Assignment

Red Hat Enterprise Linux 7 supports device assignment of NVIDIA professional series graphics devices (GRID and Quadro) as a secondary graphics device to emulated VGA.

Para-Virtualized Ticketlocks

Red Hat Enterprise Linux 7 supports para-virtualized ticketlocks (pvticketlocks) that improve performance of Red Hat Enterprise Linux 7 guest virtual machines running over Red Hat Enterprise Linux 7 hosts with oversubscribed CPUs.

Error Handling on Assigned PCIe Devices

If a PCIe device with Advanced Error Reporting (AER) encounters an error while assigned to a guest, the affected guest is brought down without impacting any other running guests or the host. The guests can be brought back up after the host driver for the device recovers from the error.

Q35 Chipset, PCI Express Bus and AHCI Bus Emulation

The Q35 machine type, required for PCI express bus support in KVM guest virtual machines, is available as a Technology Preview in Red Hat Enterprise Linux 7. An AHCI bus is only supported for inclusion with the Q35 machine type and is also available as a Technology Preview Red Hat Enterprise Linux 7.

VFIO-based PCI Device Assignment

The Virtual Function I/O (VFIO) user-space driver interface provides KVM guest virtual machines with an improved PCI device assignment solution. VFIO provides kernel-level enforcement of device isolation, improves security of device access and is compatible with features such as secure boot. VFIO replaces the KVM device assignment mechanism used in Red Hat Enterprise Linux 6.

Intel VT-d Large Pages

When using Virtual Function I/O (VFIO) device assignment with a KVM guest virtual machine on Red Hat Enterprise Linux 7, 1GB pages are used by the input/output memory management unit (IOMMU), thus reducing translation lookaside buffer (TLB) overhead for I/O operations. 2MB and 1GB page sizes are supported. The VT-d large pages feature is only supported on certain more recent Intel-based platforms.

KVM Clock Get Time Performance

In Red Hat Enterprise Linux 7 the **vsyscall** mechanism was enhanced to support fast reads of the clock from the user space for KVM guests. A guest virtual machine running Red Hat Enterprise Linux 7 on a Red Hat Enterprise Linux 7 host will see improved performance for applications that read the time of day frequently.

QCOW2 Version 3 Image Format

Red Hat Enterprise Linux 7 adds support for the QCOW2 version 3 Image Format.

Improved Live Migration Statistics

Information about live migration is now available to analyze and tune performance. Improved statistics include: total time, expected downtime, and bandwidth being used.

Live Migration Thread

The KVM live migration feature now uses its own thread. As a result, the guest performance is virtually not impacted by migration.

Hot Plugging of Character Devices and Serial Ports

Hot plugging new serial ports with new character devices is now supported in Red Hat Enterprise Linux 7.

Emulation of AMD Opteron G5

KVM is now able to emulate AMD Opteron G5 processors.

Support of New Intel Instructions on KVM Guests

KVM guests can use new instructions supported by Intel 22nm processors. These include:

- Floating-Point Fused Multiply-Add;
- 256-bit Integer vectors;
- big-endian move instruction (MOVBE) support;
- or HLE/HLE+.

VPC and VHDX File Formats

KVM in Red Hat Enterprise Linux 7 includes support for the Microsoft Virtual PC (VPC) and Microsoft Hyper-V virtual hard disk (VHDX) file formats. Note that these formats are supported in read-only mode only.

New Features in libguestfs

libguestfs is a set of tools for accessing and modifying virtual machine disk images. **libguestfs** included in Red Hat Enterprise Linux 7 includes a number of improvements, the most notable of which are the following:

- Secure Virtualization Using SELinux, or sVirt protection, ensures enhanced security against malicious and malformed disk images.
- Remote disks can be examined and modified, initially over Network Block Device (NBD).
- Disks can be hot plugged for better performance in certain applications.

WHQL-Certified virtio-win Drivers

Red Hat Enterprise Linux 7 includes Windows Hardware Quality Labs (WHQL) certified **virtio-win** drivers for the latest Microsoft Windows guests, namely Microsoft Windows 8, 8.1, 2012 and 2012 R2.

Host and Guest Panic Notification in KVM

A new **pvpanic** virtual device can be wired into the virtualization stack such that a guest panic can cause libvirt to send a notification event to management applications.

As opposed to the kdump mechanism, pvpanic does not need to reserve memory in the guest kernel. It is not needed to install any dependency packages in the guest. Also, the dumping procedure of pvpanic is host-controlled, therefore the guest only cooperates to a minimal extent.

To configure the panic mechanism, place the following snippet into the Domain XML **devices** element, by running **virsh edit** to open and edit the XML file:

```
<devices>
  <panic>
    <address type='isa' iobase='0x505' />
  </panic>
</devices>
```

After specifying the following snippet, the crashed domain's core will be dumped. If the domain is restarted, it will use the same configuration settings.

```
<on_crash>coredump-destroy</on_crash>
```

6.2. Xen

Red Hat Enterprise Linux 7 Xen HVM Guest

Users can use Red Hat Enterprise Linux 7 as a guest on the Xen environment.

6.3. Hyper-V

Red Hat Enterprise Linux 7 Hosted as a Generation 2 Virtual Machine

Red Hat Enterprise Linux 7 can be used as a generation 2 virtual machine in the Microsoft Hyper-V Server 2012 R2 host. In addition to the functions supported in the previous generation, generation 2 provides new functions on a virtual machine; for example: boot from a SCSI virtual hard disk, and UEFI firmware support.

6.4. VMware

open-vm-tools

To enhance performance and user experience when running Red Hat Enterprise Linux 7 as the guest on VMware ESX, Red Hat Enterprise Linux 7 includes the latest stable release of *open-vm-tools*.

Chapter 7. Linux Containers with Docker Format

Docker is an open source project that automates the deployment of applications inside Linux Containers, and provides the capability to package an application with its runtime dependencies into a container. It provides a Docker CLI command line tool for the lifecycle management of image-based containers. Linux containers enable rapid application deployment, simpler testing, maintenance, and troubleshooting while improving security. Using Red Hat Enterprise Linux 7 with Docker allows customers to increase staff efficiency, deploy third-party applications faster, enable a more agile development environment, and manage resources more tightly.

To quickly get up-and-running with **Docker** Containers, refer to [Get Started with Docker Containers](#).

Linux containers with Docker format are supported running on hosts with SELinux enabled. SELinux is not supported when the `/var/lib/docker` directory is located on a volume using the B-tree file system (Btrfs).

7.1. Components of Docker Containers

Docker works with the following fundamental components:

- *Container* – an application sandbox. Each container is based on an *image* that holds necessary configuration data. When you launch a container from an image, a writable layer is added on top of this image. Every time you commit a container (using the **docker commit** command), a new image layer is added to store your changes.
- *Image* – a static snapshot of the containers' configuration. Image is a read-only layer that is never modified, all changes are made in top-most writable layer, and can be saved only by creating a new image. Each image depends on one or more parent images.
- *Platform Image* – an image that has no parent. Platform images define the runtime environment, packages and utilities necessary for containerized application to run. The platform image is read-only, so any changes are reflected in the copied images stacked on top of it. See an example of such stacking in [Figure 7.1, "Image Layering Using Docker Format"](#).
- *Registry* – a repository of images. Registries are public or private repositories that contain images available for download. Some registries allow users to upload images to make them available to others.
- *Dockerfile* – a configuration file with build instructions for Docker images. Dockerfiles provide a way to automate, reuse, and share build procedures.

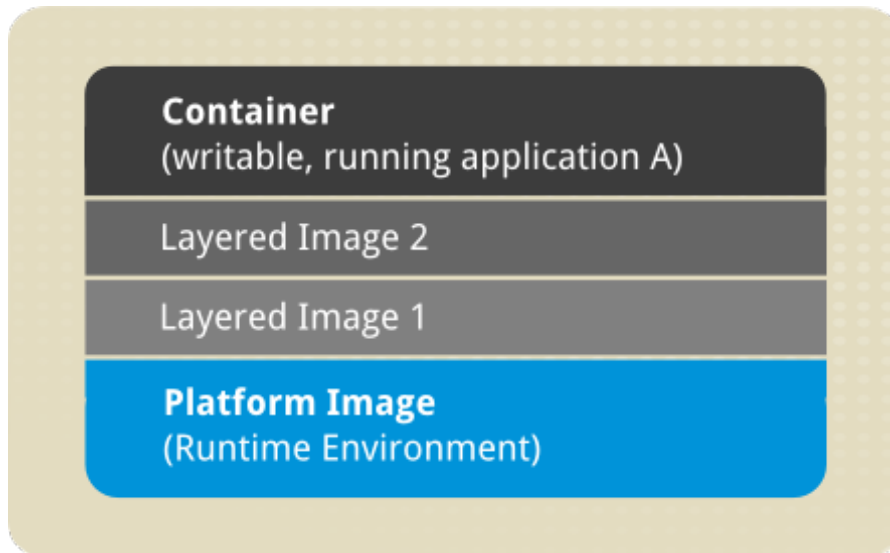


Figure 7.1. Image Layering Using Docker Format

7.2. Advantages of Using Docker

Docker brings in an API for container management, an image format and a possibility to use a remote registry for sharing containers. This scheme benefits both developers and system administrators with advantages such as:

- *Rapid application deployment* – containers include the minimal runtime requirements of the application, reducing their size and allowing them to be deployed quickly.
- *Portability across machines* – an application and all its dependencies can be bundled into a single container that is independent from the host version of Linux kernel, platform distribution, or deployment model. This container can be transferred to another machine that runs **Docker**, and executed there without compatibility issues.
- *Version control and component reuse* – you can track successive versions of a container, inspect differences, or roll-back to previous versions. Containers reuse components from the preceding layers, which makes them noticeably lightweight.
- *Sharing* – you can use a remote repository to share your container with others. Red Hat provides a registry for this purpose, and it is also possible to configure your own private repository.
- *Lightweight footprint and minimal overhead* – Docker images are typically very small, which facilitates rapid delivery and reduces the time to deploy new application containers.
- *Simplified maintenance* – Docker reduces effort and risk of problems with application dependencies.

7.3. Comparison with Virtual Machines

Virtual machines represent an entire server with all of the associated software and maintenance concerns. Docker containers provide application isolation and can be configured with minimum run-time environments. In a Docker container, the kernel and parts of the operating system infrastructure are shared. For the virtual machine, a full operating system must be included.

- You can create or destroy containers quickly and easily. Virtual Machines require full installations and require more computing resources to execute.

- Containers are lightweight, therefore, more containers than virtual machines can run simultaneously on a host machine.
- Containers share resources efficiently. Virtual machines are isolated. Therefore multiple variations of an application running in containers are also able to be very lightweight. For example, shared binaries are not duplicated on the system.
- Virtual machines can be migrated while still executing, however containers cannot be migrated while executing and must be stopped before moving from host machine to host machine.

Containers do not replace virtual machines for all use cases. Careful evaluation is still required to determine what is best for your application.

To quickly get up-and-running with **Docker** Containers, refer to [Get Started with Docker Containers](#).

The [Docker FAQ](#) contains more information about Linux Containers, Docker, subscriptions and support.

7.4. Using Docker on Red Hat Enterprise Linux 7



Note

Docker is still in development and has not yet reached version 1.0. For this reason it is not recommended to use Docker in mission-critical production environments.

Docker and Docker Registry are available via the Red Hat Enterprise Linux Extras channel. The Extras channel, which is an optional child channel of Red Hat Enterprise Linux, is intended to give customers access to select, rapidly evolving technologies. These technologies may be updated more frequently than they would otherwise be in a Red Hat Enterprise Linux minor release. The technologies delivered in the Extras channel are fully supported. Once this channel has been enabled, the packages may be installed in the usual way. For more information on installing packages or enabling channels, see [System Administrator's Guide](#).

In addition to providing installable packages for Docker and Docker Registry, Red Hat is also providing a registry of certified docker images. This registry provides pre-built solutions usable on Red Hat Enterprise Linux 7 with Docker. For more information about the registry and a list of available packages, see [Docker Images](#).

Chapter 8. System and Services

systemd

systemd is a system and service manager for Linux, and replaces SysV and Upstart used in previous releases of Red Hat Enterprise Linux. systemd is compatible with SysV and Linux Standard Base init scripts.

systemd offers, among others, the following capabilities:

- Aggressive parallelization capabilities;
- Use of socket and D-Bus activation for starting services;
- On-demand starting of daemons;
- Managing of control groups;
- Creating of system state snapshots and restoring of the system state.

For detailed information about systemd and its configuration, see [System Administrator's Guide](#).

Chapter 9. Clustering

Clusters are multiple computers (nodes) working together to increase reliability, scalability, and availability to critical production services. High Availability using Red Hat Enterprise Linux 7 can be deployed in a variety of configurations to suit varying needs for performance, high-availability, load balancing, and file sharing.

Refer to [Section 19.5, “Clustering and High Availability”](#) for a list of documents available for Red Hat Enterprise Linux 7 providing information about configuration and management of Red Hat High Availability Add-On.

9.1. Pacemaker Cluster Manager

Red Hat Enterprise Linux 7 replaces **rgmanager** with **Pacemaker** for managing cluster resources and recovering from node failures.

Some of the benefits of **Pacemaker** include:

- ▶ Automatic synchronization and versioning of the resource configuration;
- ▶ A flexible resource and fencing model that can more closely match the user's environment;
- ▶ Fencing can be used to recover from resource-level failures;
- ▶ Time-based configuration options;
- ▶ The ability to run the same resource on multiple nodes. For example, a web server or cluster file system;
- ▶ The ability to run the same resource on multiple nodes in one of two different modes. For example, a sync source and target;
- ▶ Pacemaker does not require a distributed lock manager;
- ▶ Configurable behavior when quorum is lost or multiple partitions are formed.

9.2. Keepalived and HAProxy Replace Piranha as Load Balancer

Red Hat Enterprise Linux 7 replaces the **Piranha** Load Balancer technology with **Keepalived** and **HAProxy**.

Keepalived 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 (transport layer) load balancing. **Keepalived** implements a set of checkers to dynamically and adaptively maintain and manage a load balanced server pool according to their health. **Keepalived** also implements the Virtual Router Redundancy Protocol (VRRPv2) to achieve high availability with director failover.

HAProxy is a TCP/HTTP reverse proxy which is particularly suited for high availability environments. **HAProxy** can:

- ▶ route HTTP requests depending on statically assigned cookies;
- ▶ spread the load among several servers while assuring server persistence through the use of HTTP cookies;

- ▶ switch to backup servers in the event a main server fails;
- ▶ accept connections to special ports dedicated to service monitoring;
- ▶ stop accepting connections without breaking existing ones;
- ▶ add, modify, and delete HTTP headers in both directions;
- ▶ block requests matching particular patterns;
- ▶ persist client connections to the correct application server depending on application cookies;
- ▶ report detailed status as HTML pages to authenticated users from a URI intercepted from the application.

With Red Hat Enterprise Linux 7, the Load Balancer technology is now included in the base operating system and is no longer a Red Hat Enterprise Linux Add-On.

9.3. High Availability Administration

The Pacemaker Configuration System, or **pcs**, replaces **ccs**, **ricci** and **luci** as the unified cluster configuration and administration tool. Some of the benefits of **pcs** include:

- ▶ Command-line tool;
- ▶ Ability to easily bootstrap a cluster, that is, getting the initial cluster up and running;
- ▶ Ability to configure cluster options;
- ▶ Ability to add, remove, or modify resources and their relationships to each other.

9.4. New Resource Agents

Red Hat Enterprise Linux 7 includes a number of resource agents. A resource agent is a standardized interface for a cluster resource. A resource agent translates a standard set of operations into steps specific to the resource or application, and interprets their results as success or failure.

Chapter 10. Compiler and Tools

10.1. GCC Toolchain

In Red Hat Enterprise Linux 7, the `gcc` toolchain is based on the `gcc-4.8.x` release series, and includes numerous enhancements and bug fixes relative to the Red Hat Enterprise Linux 6 equivalent. Similarly, Red Hat Enterprise Linux 7 includes `binutils-2.23.52.x`.

These versions correspond to the equivalent tools in Red Hat Developer Toolset 2.1; a detailed comparison of Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7 **gcc** and **binutils** versions can be seen here:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Developer_Toolset/2/html/User_Guide/index.html

Notable highlights of the Red Hat Enterprise Linux 7 toolchain are the following:

- Experimental support for building applications compliant with C++11 (including full C++11 language support) and some experimental support for C11 features;
- Improved support for programming parallel applications, including OpenMP v3.1, C++11 Types and GCC Built-ins for Atomic Memory Access and experimental support for transactional memory (including Intel RTM/HLE intrinsics, built-ins, and code generation);
- A new local register allocator (LRA), improving code performance;
- DWARF4 is now used as the default debug format;
- A variety of new architecture-specific options;
- Support for AMD family 15h and 16h processors;
- Link-time optimization support;
- Enhanced warnings and diagnostics;
- A variety of new Fortran features.

10.2. GLIBC

In Red Hat Enterprise Linux 7, the **glibc** libraries (**libc**, **libm**, **libpthread**, NSS plug-ins, and others) are based on the **glibc** 2.17 release, which includes numerous enhancements and bug fixes relative to the Red Hat Enterprise Linux 6 equivalent.

Notable highlights of the Red Hat Enterprise Linux 7 glibc libraries are the following:

- Experimental ISO C11 support;
- New Linux interfaces: **prlimit**, **prlimit64**, **fanotify_init**, **fanotify_mark**, **clock_adjtime**, **name_to_handle_at**, **open_by_handle_at**, **syncfs**, **setns**, **sendmmsg**, **process_vm_readv**, **process_vm_writev**;
- New optimized string functions for AMD64 and Intel 64 architectures using Streaming SIMD Extensions (SSE), Supplemental Streaming SIMD Extensions 3 (SSSE3), Streaming SIMD Extensions 4.2 (SSE4.2), and Advanced Vector Extensions (AVX);

- ▶ New optimized string functions for IBM PowerPC and IBM POWER7;
- ▶ New optimized string functions for IBM S/390 and IBM System z with specifically optimized routines for IBM System z10 and IBM zEnterprise 196;
- ▶ New locales: `os_RU`, `bem_ZA`, `en_ZA`, `ff_SN`, `sw_KE`, `sw_TZ`, `lb_LU`, `wae_CH`, `yue_HK`, `lij_IT`, `mhr_RU`, `bho_IN`, `unm_US`, `es_CU`, `ta_LK`, `ayc_PE`, `doi_IN`, `ia_FR`, `mni_IN`, `nhn_MX`, `niu_NU`, `niu_NZ`, `sat_IN`, `szl_PL`, `mag_IN`;
- ▶ New encodings: `CP770`, `CP771`, `CP772`, `CP773`, `CP774`;
- ▶ New interfaces: **`scandirat`**, **`scandirat64`**;
- ▶ Checking versions of the `FD_SET`, `FD_CLR`, `FD_ISSET`, `poll`, and `ppoll` file descriptors added;
- ▶ Caching of the netgroup database is now supported in the **`nscd`** daemon;
- ▶ The new function **`secure_getenv()`** allows secure access to the environment, returning `NULL` if running in a `SUID` or `SGID` process. This function replaces the internal function **`__secure_getenv()`**;
- ▶ The **`crypt()`** function now fails if passed salt bytes that violate the specification for those values. On Linux, the **`crypt()`** function will consult the `/proc/sys/crypto/fips_enabled` file to determine if FIPS mode is enabled, and fail on encrypted strings using the Message-Digest algorithm 5 (MD5) or Data Encryption Standard (DES) algorithm when the mode is enabled;
- ▶ The **`clock_*`** suite of functions (declared in `<time.h>`) is now available directly in the main C library. Previously it was necessary to link with `-lrt` to use these functions. This change has the effect that a single-threaded program that uses a function such as **`clock_gettime()`** (and is not linked with `-lrt`) will no longer implicitly load the `pthread`s library at runtime and so will not suffer the overheads associated with multi-thread support in other code such as the C++ runtime library;
- ▶ New header `<sys/auxv.h>` and function **`getauxval()`** allow easy access to the `AT_*` key-value pairs passed from the Linux kernel. The header also defines the `HWCAP_*` bits associated with the `AT_HWCAP` key;
- ▶ A new class of installed header has been documented for low-level platform-specific functionality. PowerPC added the first instance with a function to provide time base register access.

10.3. GDB

In Red Hat Enterprise Linux 7, the GDB debugger is based on the *`gdb-7.6.1`* release, and includes numerous enhancements and bug fixes relative to the Red Hat Enterprise Linux 6 equivalent.

This version corresponds to GDB in Red Hat Developer Toolset 2.1; a detailed comparison of Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7 GDB versions can be seen here:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Developer_Toolset/2/html/User_Guide/index.html

Notable new features of **GDB** included in Red Hat Enterprise Linux 7 are the following:

- ▶ Faster loading of symbols using the new **`.gdb_index`** section and the new **`gdb-add-index`** shell command. Note that this feature is already present in Red Hat Enterprise Linux 6.1 and later;
- ▶ **`gdbserver`** now supports standard input/output (STDIO) connections, for example: **`(gdb) target remote | ssh myhost gdbserver - hello`**;

- Improved behavior of the **watch** command using the **-location** parameter;
- Virtual method tables can be displayed by a new command, **info vtbl**;
- Control of automatic loading of files by new commands **info auto-load**, **set auto-load**, and **show auto-load**;
- Displaying absolute path to source file names using the **set filename-display absolute** command;
- Control flow recording with hardware support by a new command, **record btrace**.

Notable bug fixes in GDB included in Red Hat Enterprise Linux 7 are the following:

- The **info proc** command has been updated to work on core files;
- Breakpoints are now set on all matching locations in all inferiors;
- The file name part of breakpoint location now matches trailing components of a source file name;
- Breakpoints can now be put on inline functions;
- Parameters of the template are now put in scope when the template is instantiated.

In addition, Red Hat Enterprise Linux 7 provides a new package, *gdb-doc*, which contains the GDB Manual in PDF, HTML, and info formats. The GDB Manual was part of the main RPM package in previous versions of Red Hat Enterprise Linux.

10.4. Performance Tools

Red Hat Enterprise Linux 7 includes updates to the most recent versions of several performance tools, such as **oprofile**, **papi**, and **elfutils**, bringing performance, portability, and functionality improvements.

Moreover, Red Hat Enterprise Linux 7 premiers:

- Support for Performance Co-Pilot;
- SystemTap support for (DynInst-based) instrumentation that runs entirely in unprivileged user space, as well as efficient (Byteman-based) pinpoint probing of Java applications;
- Valgrind support for hardware transactional memory and improvements in modeling vector instructions.

10.4.1. Performance Co-Pilot

Red Hat Enterprise Linux 7 introduces support for Performance Co-Pilot (PCP), a suite of tools, services, and libraries for acquisition, archiving, and analysis of system-level performance measurements. Its light-weight, distributed architecture makes it particularly well suited to centralized analysis of complex systems.

Performance metrics can be added using the Python, Perl, C++ and C interfaces. Analysis tools can use the client APIs (Python, C++, C) directly, and rich web applications can explore all available performance data using a JSON interface.

For further information, consult the extensive man pages in the *pcp* and *pcp-libs-devel* packages. The *pcp-doc* package installs documentation in the `/usr/share/doc/pcp-doc/*` directory, which also includes these two free and open books from the upstream project:

<http://performancecopilot.org/doc/pcp-users-and-administrators-guide.pdf>

<http://performancecopilot.org/doc//pcp-programmers-guide.pdf>

10.4.2. SystemTap

Red Hat Enterprise Linux 7 includes *systemtap* version 2.4, which brings several new capabilities. These include optional pure user-space script execution, richer and more efficient Java probing, virtual machine probing, improved error messages, and a number of bug fixes and new features. In particular, the following:

- ▶ Using the **dyninst** binary-editing library, **SystemTap** can now execute some scripts purely at user-space level; no kernel or root privileges are used. This mode, selected using the **stap --dyninsti** command, enables only those types of probes or operations that affect only the user's own processes. Note that this mode is incompatible with programs that throw C++ exceptions;
- ▶ A new way of injecting probes into Java applications is supported in conjunction with the **byteman** tool. New SystemTap probe types, **java("com.app").class("class_name").method("name(signature)").***, enable probing of individual method **enter** and **exit** events in an application, without system-wide tracing;
- ▶ A new facility has been added to the SystemTap driver tooling to enable remote execution on a libvirt-managed KVM instance running on a server. It enables automated and secure transfer of a compiled SystemTap script to a virtual machine guest across a dedicated secure **virtio-serial** link. A new guest-side daemon loads the scripts and transfers their output back to the host. This way is faster and does not require IP-level networking connection between the host and the guest. To test this function, run the following command:

```
stap --remote=libvirt://MyVirtualMachine
```

- ▶ In addition, a number of improvements have been made to SystemTap's diagnostic messages:
 - Many error messages now contain cross-references to the related manual pages. These pages explain the errors and suggest corrections;
 - If a script input is suspected to contain typographic errors, a sorted suggestion list is offered to the user. This suggestion facility is used in a number of contexts when user-specified names may mismatch acceptable names, such as probed function names, markers, variables, files, aliases, and others;
 - Diagnostic duplicate-elimination has been improved;
 - ANSI coloring has been added to make messages easier to understand.

10.4.3. Valgrind

Red Hat Enterprise Linux 7 includes **Valgrind**, an instrumentation framework that includes a number of tools to profile applications. This version is based on the **Valgrind** 3.9.0 release and includes numerous improvements relative to the Red Hat Enterprise Linux 6 and Red Hat Developer Toolset 2.1 counterparts, which were based on **Valgrind** 3.8.1.

Notable new features of **Valgrind** included in Red Hat Enterprise Linux 7 are the following:

- ▶ Support for IBM System z Decimal Floating Point instructions on hosts that have the DFP facility installed;
- ▶ Support for IBM POWER8 (Power ISA 2.07) instructions;
- ▶ Support for Intel AVX2 instructions. Note that this is available only on 64-bit architectures;

- ▶ Initial support for Intel Transactional Synchronization Extensions, both Restricted Transactional Memory (RTM) and Hardware Lock Elision (HLE);
- ▶ Initial support for Hardware Transactional Memory on IBM PowerPC;
- ▶ The default size of the translation cache has been increased to 16 sectors, reflecting the fact that large applications require instrumentation and storage of huge amounts of code. For similar reasons, the number of memory mapped segments that can be tracked has been increased by a factor of 6. The maximum number of sectors in the translation cache can be controlled by the new flag **--num-transtab-sectors**;
- ▶ **Valgrind** no longer temporarily creates a mapping of the entire object to read from it. Instead, reading is done through a small fixed sized buffer. This avoids virtual memory spikes when **Valgrind** reads debugging information from large shared objects;
- ▶ The list of used suppressions (displayed when the **-v** option is specified) now shows, for each used suppression, the file name and line number where the suppression is defined;
- ▶ A new flag, **--sigill-diagnostics** can now be used to control whether a diagnostic message is printed when the just-in-time (JIT) compiler encounters an instruction it cannot translate. The actual behavior — delivery of the SIGILL signal to the application — is unchanged.
- ▶ The **Memcheck** tool has been improved with the following features:

- Improvements in handling of vector code, leading to significantly fewer false error reports. Use the **-partial-loads-ok=yes** flag to get the benefits of these changes;
- Better control over the leak checker. It is now possible to specify which kind of leaks (definite, indirect, possible, and reachable) should be displayed, which should be regarded as errors, and which should be suppressed by a given leak suppression. This is done using the options **--show-leak-kinds=kind1,kind2,...**, **--errors-for-leak-kinds=kind1,kind2,...** and an optional **match-leak-kinds:** line in suppression entries, respectively;

Note that generated leak suppressions contain this new line and are therefore more specific than in previous releases. To get the same behavior as previous releases, remove the **match-leak-kinds:** line from generated suppressions before using them;

- Reduced **possible leak** reports from the leak checker by the use of better heuristics. The available heuristics provide detection of valid interior pointers to `std::string`, to `new[]` allocated arrays with elements having destructors, and to interior pointers pointing to an inner part of a C++ object using multiple inheritance. They can be selected individually using the **--leak-check-heuristics=heur1,heur2,...** option;
 - Better control of stacktrace acquisition for heap-allocated blocks. Using the **--keep-stacktraces** option, it is possible to control independently whether a stack trace is acquired for each allocation and deallocation. This can be used to create better "use after free" errors or to decrease Valgrind's resource consumption by recording less information;
 - Better reporting of leak suppression usage. The list of suppressions used (shown when the **-v** option is specified) now shows, for each leak suppression, how many blocks and bytes it suppressed during the last leak search.
- ▶ The Valgrind GDB server integration has been improved with the following monitoring commands:
 - A new monitor command, **v.info open_fds**, that gives the list of open file descriptors and additional details;

- A new monitor command, **v.info execontext**, that shows information about the stack traces recorded by Valgrind;
- A new monitor command, **v.do expensive_sanity_check_general**, to run certain internal consistency checks.

10.5. Programming Languages

Ruby 2.0.0

Red Hat Enterprise Linux 7 provides the latest Ruby version, 2.0.0. The most notable of the changes between version 2.0.0 and 1.8.7 included in Red Hat Enterprise Linux 6 are the following:

- New interpreter, YARV (yet another Ruby VM), which significantly reduces loading times, especially for applications with large trees or files;
- New and faster "Lazy Sweep" garbage collector;
- Ruby now supports string encoding;
- Ruby now supports native threads instead of green threads.

For more information about Ruby 2.0.0, consult the upstream pages of the project: <https://www.ruby-lang.org/en/>.

Python 2.7.5

Red Hat Enterprise Linux 7 includes Python 2.7.5, which is the latest Python 2.7 series release. This version contains many improvements in performance and provides forward compatibility with Python 3. The most notable of the changes in Python 2.7.5 are the following:

- An ordered dictionary type;
- A faster I/O module;
- Dictionary comprehensions and set comprehensions;
- The sysconfig module.

For the full list of changes, see <http://docs.python.org/dev/whatsnew/2.7.html>

Java 7 and Multiple JDKs

Red Hat Enterprise Linux 7 features OpenJDK7 as the default Java Development Kit (JDK) and Java 7 as the default Java version. All Java 7 packages (*java-1.7.0-openjdk*, *java-1.7.0-oracle*, *java-1.7.1-ibm*) allow installation of multiple versions in parallel, similarly to the kernel.

The ability of parallel installation allows users to try out multiple versions of the same JDK simultaneously, to tune performance and debug problems if needed. The precise JDK is selectable through **/etc/alternatives/** as before.



Important

The Optional channel must be enabled in order to successfully install the *java-1.7.1-ibm-jdbc* or *java-1.7.1-ibm-plugin* packages from the Supplementary channel. The Optional channel contains packages that satisfy dependencies of the desired Java packages. Before installing packages from the Optional and Supplementary channels, see [Scope of Coverage Details](#). Information on subscribing to the Optional and Supplementary channels can be found in the Red Hat Knowledgebase solution [How to access Optional and Supplementary channels](#).

Chapter 11. Networking

NetworkManager

NetworkManager has been significantly enhanced to configure and monitor all the networking features for enterprise class servers and for desktop applications.

For the enterprise data centers, **NetworkManager** can be used for tasks such as basic networking configuration, network teaming, configuring virtual LANs, bridges, bonds, IPv6, VPNs, assigning interfaces to firewall zones, and others. For desktop servers it can manage wired and wireless networks and VPNs.

NetworkManager now comes with three types of interfaces:

- a robust CLI interface that allows users and scripts to interact with NetworkManager;
- NetworkManager TUI that is a text-based highlight-and-select type of interface;
- NetworkManager GUI that is more suitable for GUI desktop environments.

NetworkManager can also work side by side with initscripts if the system administrators prefer a mixed environment. NetworkManager also has full support for D-Bus as well as OpenLMI interfaces.

Networking Team Driver

In the past, the bonding driver was used for all types of link aggregation, which created various challenges. Network Teaming has been introduced as an alternative to bonding for link aggregation. The Team driver offers performance and flexibility improvements. Unlike with bonding, the control and management interface is located in user space and the fast data path is in kernel space. The Team driver supports all of the features supported by the bonding driver. A migration tool, **bond2team**, to assist with migration from bonding to teaming is also available.

Precision Time Protocol

Precision Time Protocol, or PTP, as defined in the IEEE 1588 standard, is fully supported in Red Hat Enterprise Linux 7. PTP can be used for precisely synchronizing distributed system clocks. It is capable of achieving clock accuracy in the sub-microsecond range when used in conjunction with PTP-enabled hardware devices. When used in combination with **ntpd** or **chrony**, it can be used to accurately synchronize time from the host to virtual machines. PTP also has the capability to use clock signals from GPS satellites, thus providing the same exact sub-microsecond accuracy across the globe.

chrony Suite

The **chrony** suite of utilities is available to update the system clock on systems that do not fit into the conventional permanently networked, always on, dedicated server category. The **chrony** suite should be considered for all systems which are frequently suspended or otherwise intermittently disconnected and reconnected to a network. Mobile and virtual systems for example.

Dynamic Firewall Daemon, firewalld Suite

Red Hat Enterprise Linux 7 includes the dynamic firewall daemon, **firewalld**, which provides a dynamically managed firewall with support for network "zones" to assign a level of trust to a network and its associated connections and interfaces. It has support for **IPv4** and **IPv6** firewall settings. It supports Ethernet bridges and has a separation of runtime and permanent configuration options. It also has an interface for services or applications to add firewall rules directly.

DNSSEC

DNSSEC is a set of Domain Name System Security Extensions (DNSSEC) that allows clients to determine origin authentication of **DNS** data, authenticated denial of existence and data integrity. DNSSEC prevents man-in-the-middle attacks in which active eavesdropping or intercepted communication occurs between two systems.

DDoS Protection

Distributed Denial of Service (DDoS) attacks are increasing, and becoming commonplace, as more and more products and services become dependent on delivering services over the Internet. The **SYNPROXY** module is designed to protect the system against common SYN-floods and ACK-floods, but can also be adjusted to protect against SYN-ACK floods. The **SYNPROXY** module filters out false SYN-ACK and ACK packets before the socket enters the "listen" state lock.

Support for 40 Gigabit NICs

Red Hat Enterprise Linux 7 supports 40 Gigabit network interface controllers (NICs) from multiple hardware partners. This provides support for 40 Gigabit Ethernet link speeds enabling faster network communication for applications and systems. Note that the **ethtool** utility will report interface link speeds up to 40Gb data rates.

WiGig 60 GHz Band (IEEE 802.11ad)

WiGig allows devices to wirelessly communicate at multi-gigabit speeds (up to 7 Gbps). This is nearly 50 times faster than defined in the IEEE 802.11n wireless networking standard.

Network Namespaces

Network namespaces provides a lightweight container-based virtualization that allows virtual network stacks to be associated with a process group. It creates an isolated copy of the networking data structures such as the interface list, sockets, routing table, the **/proc/net/** directory, port numbers, and so on. Network namespaces is managed through the **ip** interface (sometimes also referred to as **iproute2**), namely by the **ip netns** command.

Trusted Network Connect

Red Hat Enterprise Linux 7 introduces the Trusted Network Connect functionality as a Technology Preview. Trusted Network Connect is used with existing network access control (NAC) solutions, such as TLS, 802.1X, or IPsec to integrate end point posture assessment; that is, collecting an end point's system information (such as operating system configuration settings, installed packages, and others, termed as integrity measurements). Trusted Network Connect is used to verify these measurements against network access policies before allowing the end point to access the network.

SR-IOV Functionality in the qlcnict Driver

Support for Single-Root I/O virtualization (SR-IOV) has been added to the **qlcnict** driver as a Technology Preview. Support for this functionality will be provided directly by QLogic, and customers are encouraged to provide feedback to QLogic and Red Hat. Other functionality in the qlcnict driver remains fully supported.

FreeRADIUS 3.0.1

Red Hat Enterprise Linux 7 includes FreeRADIUS version 3.0.1, which provides a number of new features the most notable of which are the following:

- RadSec, a protocol for transporting RADIUS datagrams over TCP and TLS;
- YubiKey support;

- Connection pooling. The **radiusd** server maintains connections to a variety of back ends (SQL, LDAP, and others). Connection pooling offers greater throughput with lower resource demands;
- The syntax of the server's configuration programming language, unlang, has been expanded;
- Improved support for site-specific and vendor-specific attributes;
- Improved debugging which highlights problems in verbose output;
- SNMP trap generation;
- Improved WiMAX support;
- EAP-PWD support.

OpenLMI

Red Hat Enterprise Linux 7 features the OpenLMI project, which provides a common infrastructure for the management of Linux systems. It allows users to configure, manage, and monitor hardware, operating systems, and system services. OpenLMI is intended to simplify the task of configuring and managing production servers.

OpenLMI is designed to provide a common management interface to multiple versions of Red Hat Enterprise Linux. It builds on top of existing tools, providing an abstraction layer that hides much of the complexity of the underlying system from system administrators.

OpenLMI consists of a set of system management agents installed on a managed system, an OpenLMI controller, which manages the agents and provides an interface to them, and client applications or scripts which call the system managements agents through the OpenLMI controller.

OpenLMI allows users to:

- configure, manage, and monitor bare-metal production servers as well as virtual machine guests;
- configure, manage, and monitor local or remote systems;
- configure, manage, and monitor storage and networks;
- call system management functions from C/C++, Python, Java, or the command-line interface.

Please note that the OpenLMI software Provider is supported as a Technology Preview. The Software is fully functional, however, certain operations may consume excessive resources.

For more information about OpenLMI, see <http://www.openlmi.org>.

Chapter 12. Resource Management

Control Groups

Red Hat Enterprise Linux 7 features control groups, cgroups, which is a concept for organizing processes in a tree of named groups for the purpose of resource management. They provide a way to hierarchically group and label processes and a way to apply resource limits to these groups. In Red Hat Enterprise Linux 7, control groups are exclusively managed through **systemd**. Control groups are configured in systemd unit files and are managed with systemd's command line interface (CLI) tools.

Control groups and other resource management features are discussed in detail in [Resource Management Guide](#).

Chapter 13. Authentication and Interoperability

Improved Identity Management Cross-Realm Trusts to Active Directory

The following improvements have been implemented in cross-realm trusts to Active Directory feature of Red Hat Enterprise Linux:

- Multiple Active Directory domains are supported in the trusted forest;
- Access of users belonging to separate Active Directory domains in the trusted forest can be selectively disabled and enabled per-domain level;
- Manually defined POSIX identifiers for users and groups from trusted Active Directory domains can be used instead of automatically assigned identifiers;
- Active Directory users and groups coming from the trusted domains can be exported to legacy POSIX systems through LDAP compatibility tree;
- For Active Directory users exported through LDAP compatibility tree, authentication can be performed against Identity Management LDAP server. As a result, both Identity Management and trusted Active Directory users are accessible to legacy POSIX systems in a unified way.

Support of POSIX User and Group IDs In Active Directory

Identity Management implementation of cross-realm trusts to Active Directory supports existing POSIX user and group ID attributes in Active Directory. When explicit mappings are not defined on the Active Directory side, algorithmic mapping based on the user or group Security Identifier (SID) is applied.

Use of AD and LDAP sudo Providers

The AD provider is a back end used to connect to an Active Directory server. In Red Hat Enterprise Linux 7, using the AD sudo provider together with the LDAP provider is supported as a Technology Preview. To enable the AD sudo provider, add the **sudo_provider=ad** setting in the domain section of the **sssd.conf** file.

Support of CA-Less Installations

IPA supports installing without an embedded Certificate Authority with user-provided SSL certificates for the HTTP servers and Directory Servers. The administrator is responsible for issuing and rotating services and hosts certificates manually.

FreeIPA GUI Improvements

Red Hat Enterprise Linux 7 brings a number of improvements to FreeIPA graphical interface, from which the most notable are the following:

- All dialog windows can be confirmed by the **Enter** key even when the appropriate button or the dialog window does not have the focus;
- Loading of web UI is significantly faster because of compression of web UI assets and RPC communication;
- Drop-down lists can be controlled by keyboard.

Reclaiming IDs of Deleted Replicas

User and group ID ranges that belong to deleted replicas can be transferred to a suitable replica if one

exists. This prevents potential exhaustion of the ID space. Additionally, ID ranges can be managed manually with the **ipa-replica-manage** tool.

Re-Enrolling Clients Using Existing Keytab Files

A host that has been recreated and does not have its host entry disabled or removed can be re-enrolled using a previously backed up **keytab** file. This ensures easy re-enrolling of the IPA client system after the user rebuilds it.

Prompt for DNS

During server interactive installation, the user is asked whether to install the DNS component. Previously, the DNS feature was installed only when the **--setup-dns** option was passed to the installer, leading to users not being aware of the feature.

Enhanced SSHFP DNS Records

DNS support in Identity Management was extended with support for the RFC 6954 standard. This allows users to publish Elliptic Curve Digital Signature Algorithm (ECDSA) keys and SHA-256 hashes in SSH fingerprint (SSHFP) records.

Filtering Groups by Type

New flags, **--posix**, **--nonposix**, **--external**, can be used to filter groups by type:

- POSIX group is a group with the **posixGroup** object class;
- Non-POSIX group is a group which is not POSIX or external, which means the group does not have the **posixGroup** or **ipaExternalGroup** object class;
- External group is a group with the **ipaExternalGroup** class.

Improved Integration with the External Provisioning Systems

External provisioning systems often require extra data to correctly process hosts. A new free-form text field, **class** has been added to the host entries. This field can be used in automatic membership rules.

CRL and OCSP DNS Name in Certificate Profiles

A round-robin DNS name for the IPA Certificate Authority (CA) now points to all active IPA CA masters. The name is used for CRL and OCSP URLs in the IPA certificate profile. When any of the IPA CA masters is removed or unavailable, it does not affect the ability to check revocation status of any of the certificates issued by the IPA CA.

Certificates Search

The **cert-find** command no longer restricts users to searching certificates only by their serial number, but now also by:

- serial number range;
- subject name;
- validity period;
- revocation status;

► and issue date.

Marking Kerberos Service as Trusted for Delegation of User Keys

Individual Identity Management services can be marked to Identity Management tools as trusted for delegation. By checking the **ok_as_delegate** flag, Microsoft Windows clients can determine whether the user credentials can be forwarded or delegated to a specific server or not.

Samba 4.1.0

Red Hat Enterprise Linux 7 includes *samba* packages upgraded to the latest upstream version, which introduce several bug fixes and enhancements, the most notable of which is support for the SMB3 protocol in the server and client tools.

Additionally, SMB3 transport enables encrypted transport connections to Windows servers that support SMB3, as well as Samba servers. Also, Samba 4.1.0 adds support for server-side copy operations. Clients making use of server-side copy support, such as the latest Windows releases, should experience considerable performance improvements for file copy operations.



Warning

The updated *samba* packages remove several already deprecated configuration options. The most important are the server roles **security = share** and **security = server**. Also the web configuration tool SWAT has been completely removed. More details can be found in the Samba 4.0 and 4.1 release notes:

<https://www.samba.org/samba/history/samba-4.0.0.html>

<https://www.samba.org/samba/history/samba-4.1.0.html>

Note that several **tdb** files have been updated. This means that all **tdb** files are upgraded as soon as you start the new version of the **smbd** daemon. You cannot downgrade to an older Samba version unless you have backups of the tdb files.

For more information about these changes, refer to the Release Notes for Samba 4.0 and 4.1 mentioned above.

Chapter 14. Security

OpenSSH chroot Shell Logins

Generally, each Linux user is mapped to an SELinux user using SELinux policy, enabling Linux users to inherit the restrictions placed on SELinux users. There is a default mapping in which Linux users are mapped to the SELinux `unconfined_u` user.

In Red Hat Enterprise Linux 7, the **ChrootDirectory** option for chrooting users can be used with unconfined users without any change, but for confined users, such as `staff_u`, `user_u`, or `guest_u`, the SELinux `selinuxuser_use_ssh_chroot` variable has to be set. Administrators are advised to use the `guest_u` user for all chrooted users when using the **ChrootDirectory** option to achieve higher security.

OpenSSH - Multiple Required Authentications

Red Hat Enterprise Linux 7 supports multiple required authentications in SSH protocol version 2 using the **AuthenticationMethods** option. This option lists one or more comma-separated lists of authentication method names. Successful completion of all the methods in any list is required for authentication to complete. This enables, for example, requiring a user to have to authenticate using the public key or GSSAPI before they are offered password authentication.

GSS Proxy

GSS Proxy is the system service that establishes GSS API Kerberos context on behalf of other applications. This brings security benefits; for example, in a situation when the access to the system keytab is shared between different processes, a successful attack against that process leads to Kerberos impersonation of all other processes.

Changes in NSS

The `nss` packages have been upgraded to upstream version 3.15.2. Message-Digest algorithm 2 (MD2), MD4, and MD5 signatures are no longer accepted for online certificate status protocol (OCSP) or certificate revocation lists (CRLs), consistent with their handling for general certificate signatures.

Advanced Encryption Standard Galois Counter Mode (AES-GCM) Cipher Suite (RFC 5288 and RFC 5289) has been added for use when TLS 1.2 is negotiated. Specifically, the following cipher suites are now supported:

- ▶ `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`;
- ▶ `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`;
- ▶ `TLS_DHE_RSA_WITH_AES_128_GCM_SHA256`;
- ▶ `TLS_RSA_WITH_AES_128_GCM_SHA256`.

New Boolean Names

Several SELinux boolean names have been changed to be more domain-specific. The old names can still be used, however, only the new names will appear in the lists of booleans.

The old boolean names and their respective new names are available from the `/etc/selinux/<policy_type>/booleans.subs_dist` file.

SCAP Workbench

SCAP Workbench is a GUI front end that provides scanning functionality for SCAP content. SCAP Workbench is included as a Technology Preview in Red Hat Enterprise Linux 7.

You can find detailed information on the website of the upstream project:

<https://fedorahosted.org/scap-workbench/>

OSCAP Anaconda Add-On

Red Hat Enterprise Linux 7 introduces the OSCP Anaconda add-on as a Technology Preview. The add-on integrates OpenSCAP utilities with the installation process and enables installation of a system following restrictions given by SCAP content.

Chapter 15. Subscription Management

Red Hat Enterprise Linux 7 is available using the Red Hat Subscription Management services. The following [Knowledge Base article](#) provides a brief overview and instructions on how to register your Red Hat Enterprise Linux 7 system with Red Hat Subscription Management.

Certificate-Based Entitlements

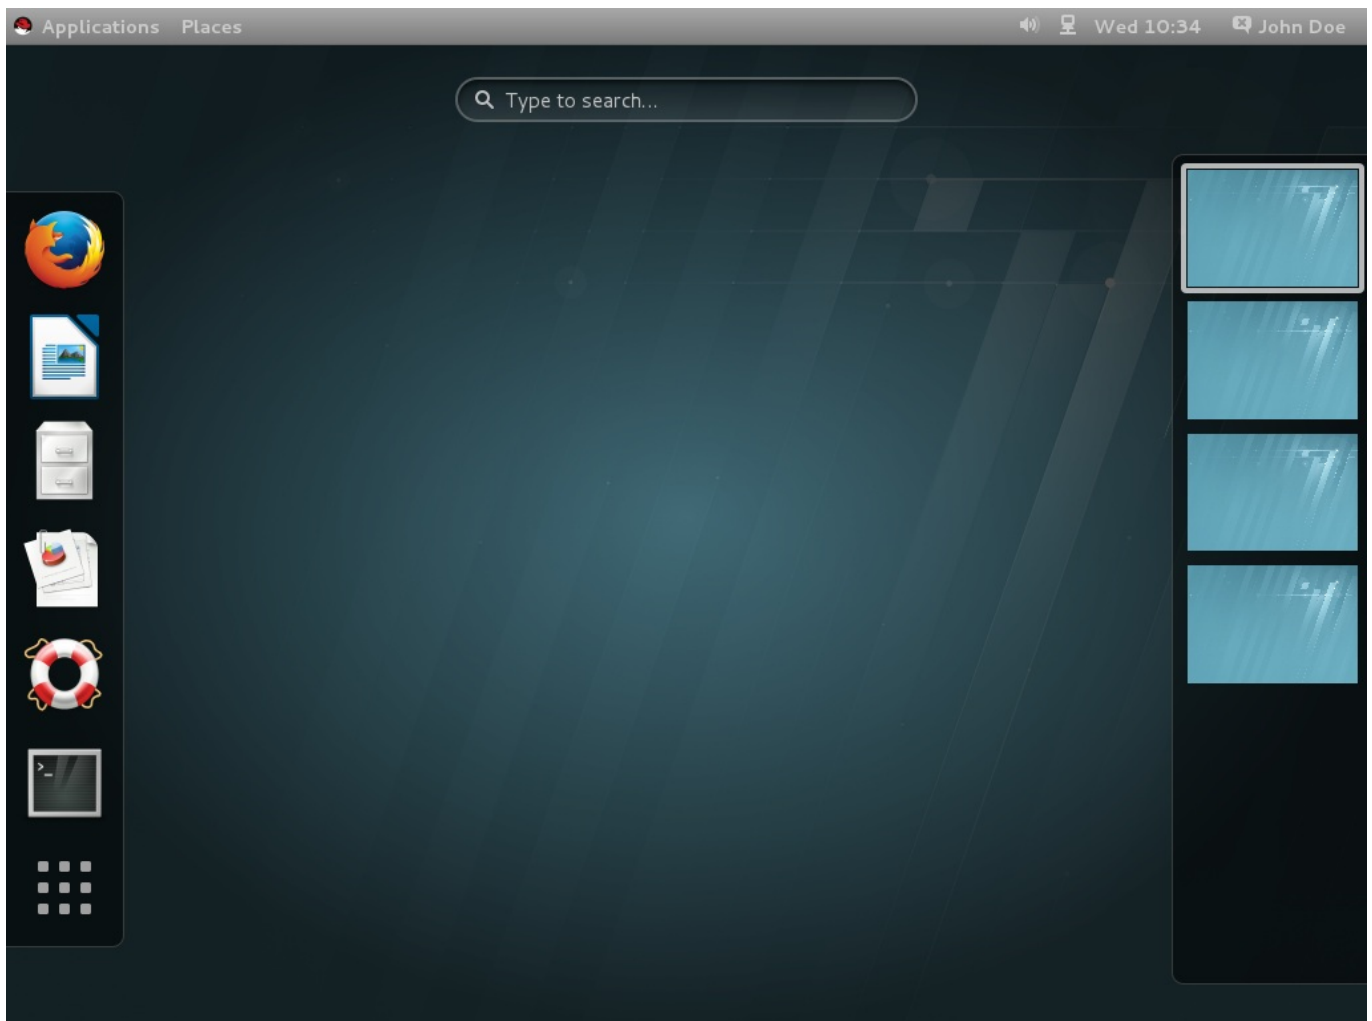
Red Hat Enterprise Linux 7 supports new certificate-based entitlements through the **subscription-manager** tool. Legacy entitlements are also supported for Satellite users to provide a transition for users using Red Hat Enterprise Linux 5 and 6. Note that registering to Red Hat Network Classic using the **rhn_register** or **rhnreg_ks** tools will not work on Red Hat Enterprise Linux 7. You can use the mentioned tools to register to Red Hat Satellite or Proxy versions 5.6 only.

Chapter 16. Desktop

16.1. GNOME 3

Red Hat Enterprise Linux 7 features the next major version of the GNOME Desktop, GNOME 3. The user experience of GNOME 3 is largely defined by GNOME Shell, which replaces the GNOME 2 desktop shell. Apart from window management, GNOME Shell provides the top bar on the screen, which hosts the "system status" area in the top right, a clock, and a hot corner that switches to **Activities Overview**, which provides easy access to applications and windows.

The default GNOME Shell interface in Red Hat Enterprise Linux 7 is GNOME Classic which features a window list at the bottom of the screen and traditional **Applications** and **Places** menus.



For more information about GNOME 3, consult GNOME help. To access it, press the **Super (Windows)** key to enter the **Activities Overview**, type **help**, and then press **Enter**.

For more information about GNOME 3 Desktop deployment, configuration and administration, see the [Desktop Migration and Administration Guide](#).

GTK+ 3

GNOME 3 uses the GTK+ 3 library which can be installed in parallel with GTK+ 2. Both GTK+ and GTK+ 3 are available in Red Hat Enterprise Linux 7. Existing GTK+ 2 applications will continue to work in GNOME 3.

GNOME Boxes

Red Hat Enterprise Linux 7 introduces a lightweight graphical desktop virtualization tool used to view and access virtual machines and remote systems. GNOME Boxes provides a way to test different operating systems and applications from the desktop with minimal configuration.

16.2. KDE

Red Hat Enterprise Linux 7 features KDE Plasma Workspaces version 4.10 and the latest version of KDE Platform and Applications. To learn more about the release, consult <http://www.kde.org/announcements/4.10/>.

KScreen

Configuration of multiple displays is improved with **KScreen**, a new screen management application for KDE. **KScreen** provides a new user interface for monitor configuration and automatic saving and restoring of profiles for connected monitors. For detailed information about KScreen, see <http://community.kde.org/Solid/Projects/ScreenManagement>.

Chapter 17. Web Servers and Services

Apache HTTP Server 2.4

Version 2.4 of the Apache HTTP Server (**httpd**) is included in Red Hat Enterprise Linux 7, and offers a range of new features:

- an enhanced version of the "Event" processing module, improving asynchronous request process and performance;
- native FastCGI support in the **mod_proxy** module;
- support for embedded scripting using the Lua language.

More information about the features and changes in **httpd** 2.4 can be found at http://httpd.apache.org/docs/2.4/new_features_2_4.html. A guide to adapting configuration files is also available: <http://httpd.apache.org/docs/2.4/upgrading.html>.

MariaDB 5.5

MariaDB is the default implementation of MySQL in Red Hat Enterprise Linux 7. MariaDB is a community-developed fork of the MySQL database project, and provides a replacement for MySQL. MariaDB preserves API and ABI compatibility with MySQL and adds several new features; for example, a non-blocking client API library, the Aria and XtraDB storage engines with enhanced performance, better server status variables, and enhanced replication.

Detailed information about MariaDB can be found at <https://mariadb.com/kb/en/what-is-mariadb-55/>.

PostgreSQL 9.2

PostgreSQL is an advanced Object-Relational database management system (DBMS). The *postgresql* packages include the PostgreSQL server package, client programs, and libraries needed to access a PostgreSQL DBMS server.

Red Hat Enterprise Linux 7 features version 9.2 of PostgreSQL. For a list of new features, bug fixes and possible incompatibilities against version 8.4 packaged in Red; Hat Enterprise; Linux; 6, please refer to the upstream release notes:

- <http://www.postgresql.org/docs/9.2/static/release-9-0.html>
- <http://www.postgresql.org/docs/9.2/static/release-9-1.html>
- <http://www.postgresql.org/docs/9.2/static/release-9-2.html>

Or the PostgreSQL wiki pages:

- http://wiki.postgresql.org/wiki/What's_new_in_PostgreSQL_9.0
- http://wiki.postgresql.org/wiki/What's_new_in_PostgreSQL_9.1
- http://wiki.postgresql.org/wiki/What's_new_in_PostgreSQL_9.2

Chapter 18. Red Hat Software Collections

Red; Hat Software Collections is a Red; Hat content set that provides a set of dynamic programming languages, database servers, and related packages that you can install and use on all supported releases of Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7.

Dynamic languages, database servers, and other tools distributed with Red; Hat Software Collections do not replace the default system tools provided with Red Hat Enterprise Linux 7, nor are they used in preference to these tools.

Red; Hat Software Collections uses an alternative packaging mechanism based on the **sc1** utility to provide a parallel set of packages. This set allows for optional use of alternative package versions on Red Hat Enterprise Linux 7. By using the **sc1** utility, users can pick and choose at any time which package version they want to run.



Important

Red; Hat Software Collections has a shorter life cycle and support term than Red; Hat Enterprise; Linux. For more information, see the [Red; Hat Software Collections Product Life Cycle](#).

See the [Red; Hat Software Collections 1.1 Release Notes](#) for important information about the Red; Hat Software Collections 1.1 release. Read this book if you want to learn about the components included in the set. This book also documents the system requirements and known problems.

See the [Red; Hat Developer Toolset 2.1 User Guide](#) for more information about installing and using Red; Hat Developer Toolset.

Chapter 19. Documentation

Documentation for Red Hat Enterprise Linux 7 is comprised of several separate documents. Each of these documents belongs to one or more of the following subject areas:

- Release Documentation;
- Installation and Deployment;
- Security
- Tools and Performance;
- Clustering
- Virtualization.

19.1. Release Documentation

Release Notes

The [Release Notes](#) document the major new features in Red Hat Enterprise Linux 7.

Migration Planning Guide

The Red Hat Enterprise Linux [Migration Planning Guide](#) documents migration from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7.

Desktop Migration and Administration Guide

The [Desktop Migration and Administration Guide](#) is a guide to the GNOME 3 Desktop migration planning, deployment, configuration, and administration on Red Hat Enterprise Linux 7.

19.2. Installation and Deployment

Installation Guide

The [Installation Guide](#) documents relevant information regarding the installation of Red Hat Enterprise Linux 7. This book also covers advanced installation methods such as kickstart, PXE installations, and installations over VNC, as well as common post-installation tasks.

System Administrator's Guide

The [System Administrator's Guide](#) provides information about deploying, configuring, and administering Red Hat Enterprise Linux 7.

Storage Administration Guide

The [Storage Administration Guide](#) provides instructions on how to effectively manage storage devices and file systems on Red Hat Enterprise Linux 7. It is intended for use by system administrators with intermediate experience in Red Hat Enterprise Linux.

Global File System 2

The [Global File System 2](#) book provides information about configuring and maintaining Red Hat GFS2 (Global File System 2) in Red Hat Enterprise Linux 7.

Logical Volume Manager Administration

The [Logical Volume Manager Administration](#) guide describes the LVM logical volume manager and provides information on running LVM in a clustered environment.

Kernel Crash Dump Guide

The [Kernel Crash Dump Guide](#) documents how to configure, test, and use the kdump crash recovery service available in Red Hat Enterprise Linux 7.

19.3. Security

Security Guide

The [Security Guide](#) is designed to assist users and administrators in learning the processes and practices of securing workstations and servers against local and remote intrusion, exploitation, and malicious activity.

SELinux User's and Administrator's Guide

The [SELinux User's and Administrator's Guide](#) covers the management and use of Security-Enhanced Linux. Note that managing confined services, which was documented in a stand-alone book in Red Hat Enterprise Linux 6, is now part of the SELinux User's and Administrator's Guide.

19.4. Tools and Performance

Resource Management and Linux Containers Guide

The [Resource Management and Linux Containers Guide](#) documents tools and techniques for managing system resources and deploying LXC application containers on Red Hat Enterprise Linux 7.

Performance Tuning Guide

The [Performance Tuning Guide](#) documents how to optimize subsystem throughput in Red Hat Enterprise Linux 7.

Developer Guide

The [Developer Guide](#) describes the different features and utilities that make Red Hat Enterprise Linux 7 an ideal enterprise platform for application development.

SystemTap Beginners Guide

The [SystemTap Beginners Guide](#) provides basic instructions on how to use SystemTap to monitor different subsystems of Red Hat Enterprise Linux in finer detail.

SystemTap Reference

The [SystemTap Tapset Reference](#) guide describes the most common tapset definitions users can apply to SystemTap scripts.

19.5. Clustering and High Availability

High Availability Add-On Administration

The [High Availability Add-On Administration](#) guide provides information on how to configure and administer the High Availability Add-On in Red Hat Enterprise Linux 7.

High Availability Add-On Overview

The [High Availability Add-On Overview](#) document provides an overview of the High Availability Add-On for Red Hat Enterprise Linux 7.

High Availability Add-On Reference

[High Availability Add-On Reference](#) is a reference guide to the High Availability Add-On for Red Hat Enterprise Linux 7.

Load Balancer Administration

[Load Balancer Administration](#) is a guide to configuring and administering high-performance load balancing in Red Hat Enterprise Linux 7.

DM Multipath

The [DM Multipath](#) book guides users through configuring and administering the Device-Mapper Multipath feature for Red Hat Enterprise Linux 7.

19.6. Virtualization

Virtualization Getting Started Guide

The [Virtualization Getting Started Guide](#) is an introduction to virtualization on Red Hat Enterprise Linux 7.

Virtualization Deployment and Administration Guide

The [Virtualization Deployment and Administration Guide](#) provides information on installing, configuring, and managing virtualization on Red Hat Enterprise Linux 7.

Virtualization Security Guide

The [Virtualization Security Guide](#) provides an overview of virtualization security technologies provided by Red Hat, and provides recommendations for securing virtualization hosts, guests, and shared infrastructure and resources in virtualized environments.

Virtualization Tuning and Optimization Guide

The [Virtualization Tuning and Optimization Guide](#) covers KVM and virtualization performance. Within this guide you can find tips and suggestions for making full use of KVM performance features and options for your host systems and virtualized guests.

Chapter 20. Internationalization

20.1. Red Hat Enterprise Linux 7 International Languages

Red Hat Enterprise Linux 7 supports the installation of multiple languages and the changing of languages based on your requirements.

The following languages are supported in Red Hat Enterprise Linux 7:

- East Asian Languages - Japanese, Korean, Simplified Chinese, and Traditional Chinese;
- European Languages - English, German, Spanish, French, Italian, Portuguese Brazilian, and Russian.
- Indic Languages - Assamese, Bengali, Gujarati, Hindi, Kannada, Malayalam, Marathi, Odia, Punjabi, Tamil, and Telugu.

The table below summarizes the currently supported languages, their locales, default fonts installed, and packages required for some of the supported languages.

For more information on font configuration, see [Desktop Migration and Administration Guide](#).

Table 20.1. Language Support Matrix

| Territory | Language | Locale | Default Font (Font Package) | Input Methods |
|-----------|--------------------|-------------|--|------------------------------------|
| Brazil | Portuguese | pt_BR.UTF-8 | DejaVu Sans (dejavu-sans-fonts) | |
| France | French | fr_FR.UTF-8 | DejaVu Sans (dejavu-sans-fonts) | |
| Germany | German | de_DE.UTF-8 | DejaVu Sans (dejavu-sans-fonts) | |
| Italy | Italian | it_IT.UTF-8 | DejaVu Sans (dejavu-sans-fonts) | |
| Russia | Russian | ru_RU.UTF-8 | DejaVu Sans (dejavu-sans-fonts) | |
| Spain | Spanish | es_ES.UTF-8 | DejaVu Sans (dejavu-sans-fonts) | |
| USA | English | en_US.UTF-8 | DejaVu Sans (dejavu-sans-fonts) | |
| China | Simplified Chinese | zh_CN.UTF-8 | WenQuanYi Zen Hei Sharp (wqy-zenhei-fonts) | ibus-libpinyin, ibus-table-chinese |
| Japan | Japanese | ja_JP.UTF-8 | VL PGothic (vlgothic-p-fonts) | ibus-kkc |

| Territory | Language | Locale | Default Font (Font Package) | Input Methods |
|-----------|---------------------|-------------|---------------------------------------|----------------------------------|
| Korea | Korean | ko_KR.UTF-8 | NanumGothic (nhn-nanum-gothic-fonts) | ibus-hangul |
| Taiwan | Traditional Chinese | zh_TW.UTF-8 | AR PL UMing TW (cjkluni-uming-fonts) | ibus-chewing, ibus-table-chinese |
| India | Assamese | as_IN.UTF-8 | Lohit Assamese (lohit-assamese-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Bengali | bn_IN.UTF-8 | Lohit Bengali (lohit-bengali-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Gujarati | gu_IN.UTF-8 | Lohit Gujarati (lohit-gujarati-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Hindi | hi_IN.UTF-8 | Lohit Hindi (lohit-devanagari-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Kannada | kn_IN.UTF-8 | Lohit Kannada (lohit-kannada-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Malayalam | ml_IN.UTF-8 | Meera (smc-meera-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Marathi | mr_IN.UTF-8 | Lohit Marathi (lohit-marathi-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Odia | or_IN.UTF-8 | Lohit Oriya (lohit-oriya-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Punjabi | pa_IN.UTF-8 | Lohit Punjabi (lohit-punjabi-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Tamil | ta_IN.UTF-8 | Lohit Tamil (lohit-tamil-fonts) | ibus-m17n, m17n-db, m17n-contrib |
| | Telugu | te_IN.UTF-8 | Lohit Telugu (lohit-telugu-fonts) | ibus-m17n, m17n-db, m17n-contrib |

20.2. General Changes In Internationalization

New *yum-langpacks* Plug-In

A new Yum plug-in, *yum-langpacks* enables users to install translation subpackages for various packages for the current language locale. This plug-in also provides Yum commands that show available languages support, display the list of installed languages, allow users to install new languages, remove installed languages and also show which packages will be installed when the user wants to install new language support.

These changes can be illustrated by the following example:

To install language packs for the Marathi or Czech languages in Red Hat Enterprise Linux 6, run:

```
]# yum groupinstall marathi-support
~]# yum groupinstall czech-support
```

To install language packs for the Marathi or Czech languages in Red Hat Enterprise Linux 7, run:

```
~]# yum langinstall mr
~]# yum langinstall cs
```

Please refer to the **yum-langpacks(8)** man page for more information.

Changing Locale and Keyboard Layout Settings

localectl is a new utility used to query and change the system locale and keyboard layout settings; the settings are used in text consoles and inherited by desktop environments. **localectl** also accepts a **hostname** argument to administer remote systems over SSH.

20.3. Input Methods

Changes in IBus

Red Hat Enterprise Linux 7 includes support for the Intelligent Input Bus (IBus) version 1.5. Support for IBus is now integrated in GNOME.

- ▶ Input methods can be added using the **gnome-control-center region** command, and the **gnome-control-center keyboard** command can be used to set input hotkeys.
- ▶ For non-GNOME sessions, *ibus* can configure both XKB layouts and input methods in the **ibus-setup** tool and switch them with a hotkey.
- ▶ The default hotkey is **Super+space**, replacing **Control+space** in *ibus* included in Red Hat Enterprise Linux 6. This provides a similar UI which the user can see with the **Alt+Tab** combination. Multiple input methods can be switched using the **Alt+Tab** combination.

Predictive Input Method for IBus

ibus-typing-booster is a predictive input method for the ibus platform. It predicts complete words based on partial input. Users can select the desired word from a list of suggestions and improve their typing speed and spelling. *ibus-typing-booster* works also with the Hunspell dictionaries and can make suggestions for a language using a Hunspell dictionary.

Note that the *ibus-typing-booster* package is an optional package, and therefore will not be installed as part of the *input-methods* group by default.

For more detailed changes in input methods, see [Desktop Migration and Administration Guide](#).

20.4. Fonts

fonts-tweak-tool

A new tool, **fonts-tweak-tool** enables users to configure the default fonts per language.

20.5. Language-Specific Changes

Arabic

New Arabic fonts from Paktype are available in Red Hat Enterprise Linux 7: `paktype-ajrak`, `paktype-basic-naskh-farsi`, `paktype-basic-naskh-sindhi`, `paktype-basic-naskh-urdu`, and `paktype-basic-naskh-sa`.

Chinese

- ▶ The WQY Zenhei font is now the default font for Simplified Chinese.
- ▶ The default engine for Simplified Chinese has been changed to `ibus-libpinyin` from `ibus-pinyin` that Red Hat Enterprise Linux 6 uses.

Indic

- ▶ The new Lohit Devanagari font replaces the previous separate Lohit fonts for Hindi, Kashmiri, Konkani, Maithili, Marathi, and Nepali. Any distinct glyphs for these languages needed in the future can be handled in Lohit Devanagari with the Open Type Font `locl` tags.
- ▶ New font packages *gubbi-fonts* and *navilu-fonts* have been added for Kannada language.

Japanese

- ▶ IPA fonts are no longer installed by default
- ▶ `ibus-kkc`, the Kana Kanji Conversion, is the new default Japanese input method engine using the new `libkkc` back end. It replaces `ibus-anthy`, `anthy`, and `kasumi`.

Korean

The Nanum font is used by default now.

New Locales

Red Hat Enterprise Linux 7 supports new locales, Konkani (`kok_IN`) and Pushto (`ps_AF`).

Chapter 21. Supportability and Maintenance

ABRT 2.1

Red Hat Enterprise Linux 7 includes the **Automatic Bug Reporting Tool (ABRT)** 2.1, which features an improved user interface and the ability to send *μReports*, lightweight anonymous problem reports suitable for machine processing, such as gathering crash statistics. The set of supported languages, for which **ABRT** is capable of detecting problems, has been extended with the addition of Java and Ruby in **ABRT 2.1**.

In order to use **ABRT**, ensure that the *abrt-desktop* or the *abrt-cli* package is installed on your system. The *abrt-desktop* package provides a graphical user interface for **ABRT**, and the *abrt-cli* package contains a tool for using **ABRT** on the command line. You can also install both.

To install the package containing the graphical user interface for **ABRT**, run the following command as the **root** user:

```
~]# yum install abrt-desktop
```

To install the package that provides the command line **ABRT** tool, use the following command:

```
~]# yum install abrt-cli
```

Note that while both of the above commands cause the main **ABRT** system to be installed, you may need to install additional packages to obtain support for detecting crashes in software programmed using various languages. See the *Automatic Bug Reporting Tool (ABRT)* chapter of the [Red Hat Enterprise Linux 7 System Administrator's Guide](#) for information on additional packages available with the **ABRT** system.

Upon installation, the **abrt-d** daemon, which is the core of the **ABRT** crash-detecting service, is configured to start at boot time. You can use the following command to verify its current status:

```
~]$ systemctl is-active abrt-d.service
active
```

In order to discover as many software bugs as possible, administrators should configure **ABRT** to automatically send reports of application crashes to Red Hat. To enable the autoreporting feature, issue the following command as **root**:

```
~]# abrt-auto-reporting enabled
```

Additional Information on ABRT

- [Red Hat Enterprise Linux 7 System Administrator's Guide](#) — The *Automatic Bug Reporting Tool (ABRT)* chapter of the *Administrator's Guide* for Red Hat Enterprise Linux 7 contains detailed information on installing, configuring, and using the **ABRT** service.

Part II. Known Issues

This part describes known issues in Red Hat Enterprise Linux 7.

Chapter 22. Installation

yaboot component, BZ#1032149

Due to a bug in the **yaboot** boot loader, upgrading from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7 can fail on 64-bit PowerPC systems with a **Unknown or corrupt filesystem** error.

anaconda component, BZ#1083994

Under certain circumstances, searching for a Fibre Channel over Ethernet (FCoE) device causes a traceback error. To work around this problem, activate the required network device before entering the Storage spoke for the first time.

kernel component, BZ#1032048

Upgrading Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7 on IBM System z using the **rhelup** program causes the following error message to appear:

```
Error: Config file '/etc/zipl.conf': Line 9: section name '3.10.0-52.el7.s390x' already specified
```

This is because a kernel rescue image is installed during the upgrade using the same zipl boot label as the new kernel image. Consequently, the **zipl** program does not update the boot loader because these duplicated entries exist in the **/etc/zipl.conf** file.

To work around this problem, edit the **/etc/zipl.conf** file and rename the label of the entry pointing to the kernel rescue image. Save the configuration file and run the **zipl** program again.

anaconda component, BZ#1027737

It is not possible to rescue a system by using an iSCSI disk; when starting **anaconda** in rescue mode on a system with an iSCSI disk, **anaconda** does not allow the user to connect to the disk.

dracut component, BZ#1096979

The **DASD=** option in the CMS configuration file defines the Direct Access Storage Device (DASD) or a range of DASDs to configure during the installation. To indicate that no DASD is present, remove the option from the configuration file because specifying **DASD=none** is not valid.

anaconda component, BZ#1070104

It is not possible to install Red Hat Enterprise Linux 7 from the installation DVD using the Hardware Management Console (HMC) on IBM System z architecture.

anaconda component, BZ#1073982

Direct Access Storage Devices (DASDs) that are formatted using the **cpfmtxa** utility are not used during the kickstart installation, even if the **zerombr** command is present in the kickstart file. To work around this problem, use the **dasdfmt** utility in the kickstart **%pre** section to format DASDs, for example:

```
%pre
dasdfmt -y -d cdl -b 4096 /dev/disk/by-path/ccw-0.0.3727
%end
```

dracut component, BZ#1094773

The **SUBCHANNELS=** variable in the CMS configuration file provides required device bus IDs for the various network interfaces. On IBM System z architecture it is necessary to specify the IDs in lowercase, otherwise the installation program fails to configure the interfaces.

kernel component, BZ#1055814

When installing Red Hat Enterprise Linux 7 on UEFI-based systems, the Anaconda installer terminates unexpectedly with the following error:

```
BootLoaderError: failed to remove old efi boot entry
```

To work around this problem, edit the **Install Red Hat Enterprise Linux 7** option in the boot menu by pressing the **e** key and append the **efi_no_storage_paranoia** kernel parameter to the end of the line that begins with **linuxefi**. Then press the **F10** key to boot the modified option and start installation.

anaconda component, BZ#[1072619](#)

It is not possible to use read-only disks as hard drive installation repository sources. When specifying the **inst.repo=hd:device:path** option ensure that *device* is writable.

kernel component, BZ#1067292, BZ#1008348

Various platforms include BIOS or UEFI-assisted software RAID provided by LSI. This hardware requires the closed-source **megasr** driver, which is not included in Red Hat Enterprise Linux. Thus, platforms and adapters that depend on **megasr** are not supported by Red Hat. Also, the use of certain open-source RAID alternatives, such as the **dmraid** Disk Data Format 1 (DDF1) capability, is not currently supported on these systems

However, on certain systems, such as IBM System x servers with the ServeRAID adapter, it is possible to disable the BIOS RAID function. To do this, enter to the UEFI menu and navigate through the **System Settings** and **Devices and I/O Ports** submenus to the **Configure the onboard SCU** submenu. Then change the SCU setting from **RAID** to **nonRAID**. Save your changes and reboot the system. In this mode, the storage is configured using an open-source non-RAID LSI driver shipped with Red Hat Enterprise Linux, such as **mptsas**, **mpt2sas**, or **mpt3sas**.

To obtain the **megasr** driver for IBM systems, refer to the [IBM support page](#).

Certain Cisco Unified Computing System (UCS) platforms are also impacted by this restriction. However, it is not possible to disable the BIOS RAID function on these systems. To obtain the **megasr** driver, refer to the [Cisco support page](#).



Note

The described restriction does not apply to LSI adapters that use the **megaraid** driver. Those adapters implement the RAID functions in the adapter firmware.

grub2 component, BZ#948213

Sometimes, when using the **help** command on the GRUB 2 command line (for example, **help set** or **help ls**), the command can become unresponsive or even cause the machine to reboot.

grub2 component, BZ#824041

If Fedora 17 or later is installed on the system, the kernel uses a shared **/boot** partition. If the user then installs Red Hat Enterprise Linux 7 which uses the same **/boot** partition, the latest kernel will be used by GRUB 2. That is, Red Hat Enterprise Linux 7 can use a Fedora kernel if this is later, resulting in certain modules to be limited or non-functioning (for example, sound or networking). To work around the problem, the correct kernel must be chosen manually.

kernel component, BZ#833561

On certain Intel systems, the following error message can appear on boot:

```
[ 17.624504] ioapic: probe of 0000:00:05.4 failed with error -22
[ 17.631700] ioapic: probe of 0000:80:05.4 failed with error -22
```

This message is harmless and does not affect the user.

anaconda component, BZ#978266

When Btrfs volumes are allocated during automatic partitioning, there is not enough space to hold the swap partition as swap is a standard partition and cannot be split onto multiple disks. Consequently, an attempt to create Btrfs partitioning automatically on all direct access storage devices (DASDs) at once fails with the "not enough free space on disks" error message.

anaconda component, BZ#980483

During installation of Red Hat Enterprise Linux 7 from a DVD ISO image stored on a disk, an exception is raised when the user tries to select a different ISO image. Consequently, the installation fails with a traceback. To work around this problem, do not try to modify installation source when the hard disk source has been specified on the kernel command line.

anaconda component, BZ#959866

Under some circumstances, installing Red Hat Enterprise Linux 7 on the unified extensible firmware interface (UEFI) causes the **anaconda** utility to report the following error message:

```
BootLoaderError: failed to remove old efi boot entry
```

anaconda component, BZ#1035201

Installing Red Hat Enterprise Linux 7 on IBM System z fails if encryption and auto-partitioning is used. To work around this problem, the user can use the custom partitioning screen to create desired partitioning layout and encrypt any mount points except for the **/boot** mount point. As a result, the partitioning is accepted by the **anaconda** utility and the installation finishes successfully.

dracut component, BZ#1023039

A system with an encrypted partition does not always ask for the password to unlock the encrypted partition during boot, which causes the boot process to fail. To work around this problem, reboot the system.

anaconda component, BZ#1036128

When auto-partitioning with standard partitions is used while installing Red Hat Enterprise Linux 7 in the **virt-manager** utility, the installation fails with a "format create" error message.

anaconda component, BZ#965985

When booting in rescue mode on IBM System z architecture, the second and third rescue screens in the rescue shell are incomplete and not displayed properly.

anaconda component, BZ#873135

On IBM System z machines, if the `/boot` sector is a logical volume that extends across more than one physical volume, and the system uses direct access storage devices (DASDs) of differing sizes, installation of bootloader fails. To work around this problem, keep `/boot` small enough so that it does not extend beyond a single physical volume. Alternatively, do not use LVM for `/boot`.

python-blivet component, BZ#1075671

If the system is installed from a removable medium, and an install or storage option is selected but then changed, the installation fails. Consequently, it is not possible to reconsider decisions made during the installation process. To work around this problem, avoid accepting and then changing storage decisions if using removable install media.

anaconda component, BZ#1058858

On systems that use the `zipl` bootloader, the `>bootloader --boot-drive` value from the anaconda kickstart file is ignored. Consequently, the boot drive is chosen by `zipl` regardless of the kickstart configuration. As a result, there may be a discrepancy between the `--boot-drive` value in kickstart file and the actual drive on which the bootloader is located. This problem applies for both manually and automatically created kickstart files.

anaconda component, BZ#885011

On IBM System z machines, if the `zipl.conf` configuration file contains a kernel parameter line that exceeds 896 bytes, installation of bootloader fails due to a hardware limitation of the byte length. Consequently, the **Anaconda** installer terminates with an error message indicating a problem in the bootloader installation. To work around this problem, keep the kernel parameter list's size to a minimum. For example, choose short names for volume groups when using the Logical Volume Manager (LVM) or specify a range of direct access storage devices (DASDs) instead of listing them individually.

anaconda component, BZ#1085310

Network devices are not automatically enabled during installation unless the installation method requires network connectivity. As a consequence, a traceback error can occur during Kickstart installation due to inactive network devices. To work around this problem, set the `ksdevice=link` option on boot or add the `--device=link` option to the `ks.cfg` file to enable network devices with active links during Kickstart installation.

grub2 component, BZ#1065360

If the `kernel-debug` package is installed, it is made the default boot entry in the **GRUB** boot loader, which is unexpected behavior. To work around this problem, do not install the `kernel-debug` package as a part of the **Anaconda** installation, but install `kernel-debug` using the `yum` utility after the system is set up. As a result, the default boot entry boots the default `kernel` as expected.

yum component, BZ#1058297

Under certain rare circumstances, the **anaconda** installer does not interact correctly with **yum** and returns an error with no exception set. Since this problem occurs rarely, reattempt the installation to work around this problem. Alternatively, use the text mode installation, where this bug does not occur.

anaconda component, BZ#1087774

The source code does not handle booting on a **bnx2i** iSCSI driver correctly. As a consequence, when installing Red Hat Enterprise Linux 7, the server does not reboot automatically after the installation is completed. No workaround is currently available.

anaconda component, BZ#1085325

The **anaconda** installer does not correctly handle adding of FCoE disks. As a consequence, adding FCoE disks on the **anaconda** advance storage page fails with the following error message:

No Fibre Channel Forwarders or VN2VN Responders Found

To work around this problem, simply repeat the steps to add the FCoE disks: the configuration process produces the correct outcome when repeated. Alternatively, run the **lldpad -d** command in the **anaconda** shell before adding the FCoE disks in the **anaconda** user interface to avoid the described problem.

Chapter 23. Storage

e2fsprogs component

The **e4defrag** utility in the *e2fsprogs* package is not supported in Red Hat Enterprise Linux 7.0, and is scheduled to be removed in Red Hat Enterprise Linux 7.1.

xfsprogs component

If XFS metadata checksums are enabled specifying the **-m crc=1** option to the **mkfs.xfs** command, the kernel prints the following warning message when the file system is mounted:

Version 5 superblock detected. This kernel has EXPERIMENTAL support enabled! Use of these features in this kernel is at your own risk!

Note that this CRC functionality is available for testing in Red Hat Enterprise Linux 7.0 and is planned to be fully supported in Red Hat Enterprise Linux 7.1.

cryptsetup component, BZ#883941

When the **systemd** daemon parses the crypttab file in case of a non-LUKS device for swap with a random key, it uses the **ripemd160** hash by default, which is not allowed in FIPS mode. To work around this problem, add the **hash=** setting with an algorithm approved for FIPS to the particular crypttab line. For example: **swap /dev/sda7 /dev/urandom swap,hash=sha1**.

lvm2 component, BZ#1083633

If the kernel thin provisioning code detects a device failure or runs out of metadata space, it sets a flag on device to indicate that it needs to be checked. Currently, LVM tools do not perform this check automatically. To work around this problem, execute the **thin_check --clear-needs-check-flag** command to perform the check and remove the flag. Then run the **thin_repair** command if necessary. Alternatively, you can add **--clear-needs-check-flag** to **thin_check_options** in the global section of the **/etc/lvm.conf** configuration file to run the check automatically.

device-mapper-multipath component, BZ#1066264

In Red Hat Enterprise Linux 7, the behavior of the **kpartx** utility has been changed, so it no longer adds the letter **p** as a delimiter between a device name and a partition suffix unless the device name ends with a digit. When a device is renamed using the **multipath** utility, the previous device name is simply replaced by a new name while the suffix stays unchanged, regardless of the delimiter. Consequently, the **kpartx** behavior is not followed, and changing device names ending in a digit to names ending in a letter, or the other way round, works incorrectly when using **multipath** to rename devices. To work around this problem, choose one of the following three options:

- Remove the **multipathd** daemon and add the devices again.
- Remove the devices manually by using the **kpartx -d** command and then add them by running the **kpartx -a** command.
- Rename the devices by using the **kpartx -p p** command for device names that are supposed to contain the delimiter and they do not, and the **kpartx -p ""** command in cases when the delimiter is used redundantly.

snapper component, BZ#[1071973](#)

The **empty-pre-post** cleanup algorithm, which is used for deleting pre-post pairs of file system snapshots with empty diffs, does not work in Red Hat Enterprise Linux 7. To work around this problem, remove empty pre-post snapshot couples manually by using the **delete** command.

kernel component, BZ#1084859

The bigalloc feature for the ext4 file systems, which enables ext4 to use clustered allocation, is not supported in Red Hat Enterprise Linux 7.

Chapter 24. Kernel

kernel component, BZ#1019091

The following RAID controller cards are no longer supported. However, the **aacraid** driver still detects them. Thus, they are marked as not supported in the **dmesg** output.

- ▶ PERC 2/Si (Iguana/PERC2Si)
- ▶ PERC 3/Di (Opal/PERC3Di)
- ▶ PERC 3/Si (SlimFast/PERC3Si)
- ▶ PERC 3/Di (Iguana FlipChip/PERC3DiF)
- ▶ PERC 3/Di (Viper/PERC3DiV)
- ▶ PERC 3/Di (Lexus/PERC3DiL)
- ▶ PERC 3/Di (Jaguar/PERC3DiJ)
- ▶ PERC 3/Di (Dagger/PERC3DiD)
- ▶ PERC 3/Di (Boxster/PERC3DiB)
- ▶ Adaptec 2120S (Crusader)
- ▶ Adaptec 2200S (Vulcan)
- ▶ Adaptec 2200S (Vulcan-2m)
- ▶ Legend S220 (Legend Crusader)
- ▶ Legend S230 (Legend Vulcan)
- ▶ Adaptec 3230S (Harrier)
- ▶ Adaptec 3240S (Tornado)
- ▶ ASR-2020ZCR SCSI PCI-X ZCR (Skyhawk)
- ▶ ASR-2025ZCR SCSI SO-DIMM PCI-X ZCR (Terminator)
- ▶ ASR-2230S + ASR-2230SLP PCI-X (Lancer)
- ▶ ASR-2130S (Lancer)
- ▶ AAR-2820SA (Intruder)
- ▶ AAR-2620SA (Intruder)
- ▶ AAR-2420SA (Intruder)
- ▶ ICP9024RO (Lancer)
- ▶ ICP9014RO (Lancer)
- ▶ ICP9047MA (Lancer)
- ▶ ICP9087MA (Lancer)
- ▶ ICP5445AU (Hurricane44)

- ICP9085LI (Marauder-X)
- ICP5085BR (Marauder-E)
- ICP9067MA (Intruder-6)
- Themisto Jupiter Platform
- Callisto Jupiter Platform
- ASR-2020SA SATA PCI-X ZCR (Skyhawk)
- ASR-2025SA SATA SO-DIMM PCI-X ZCR (Terminator)
- AAR-2410SA PCI SATA 4ch (Jaguar II)
- CERC SATA RAID 2 PCI SATA 6ch (DellCorsair)
- AAR-2810SA PCI SATA 8ch (Corsair-8)
- AAR-21610SA PCI SATA 16ch (Corsair-16)
- ESD SO-DIMM PCI-X SATA ZCR (Prowler)
- AAR-2610SA PCI SATA 6ch
- ASR-2240S (SabreExpress)
- ASR-4005
- ASR-4800SAS (Marauder-X)
- ASR-4805SAS (Marauder-E)
- ASR-3800 (Hurricane44)
- Adaptec 5400S (Mustang)
- Dell PERC2/QC
- HP NetRAID-4M

The following cards detected by **aacraid** are also no longer supported but they are *not* identified as not supported in the **dmesg** output:

- IBM 8i (AvonPark)
- IBM 8i (AvonPark Lite)
- IBM 8k/8k-l8 (Aurora)
- IBM 8k/8k-l4 (Aurora Lite)



Warning

Note that the **Kdump** mechanism might not work properly on the aforementioned RAID controllers.

kernel component, BZ#1061210

When the **hpsa_allow_any** option is used, the **hpsa** driver allows the use of PCI IDs that are not listed in the driver's pci-id table. Thus, cards detected when this option is used, are not supported in Red Hat Enterprise Linux 7.

kernel component, BZ#975791

The following **cciss** controllers are no longer supported:

- ▶ Smart Array 5300
- ▶ Smart Array 5i
- ▶ Smart Array 532
- ▶ Smart Array 5312
- ▶ Smart Array 641
- ▶ Smart Array 642
- ▶ Smart Array 6400
- ▶ Smart Array 6400 EM
- ▶ Smart Array 6i
- ▶ Smart Array P600
- ▶ Smart Array P800
- ▶ Smart Array P400
- ▶ Smart Array P400i
- ▶ Smart Array E200i
- ▶ Smart Array E200
- ▶ Smart Array E500
- ▶ Smart Array P700M

kernel component, BZ#1055089

The **systemd** service does not spawn the **agetty** tool on the **/dev/hvc0/ virtio console** if the **virtio console** driver is not found before loading kernel modules at system startup. As a consequence, a **TTY** terminal does not start automatically after the system boot when the system is running as a KVM guest. To work around this problem, start a **getty** on **/dev/hvc0/** after the system boot. The ISA serial device, which is used more commonly, works as expected.

kernel component, BZ#1060565

Previously applied patch is causing a memory leak when creating symbolic links over NFS. Consequently, if creating a very large number of symbolic links, in scale of hundreds of thousands, the system may report the out of memory status.

kernel component, BZ#1097468

The Linux **kernel** Non-Uniform Memory Access (NUMA) balancing does not always work correctly in Red Hat Enterprise Linux 7. As a consequence, when the **numa_balancing** parameter is set, some of the memory can move to an arbitrary non-destination node before moving to the constrained nodes, and the memory on the destination node also decreases under certain circumstances. There is currently no known workaround available.

kernel component, BZ#915855

The QLogic 1G iSCSI Adapter present in the system can cause a call trace error when the **qla4xx** driver is sharing the interrupt line with the USB sub-system. This error has no impact on the system functionality. The error can be found in the kernel log messages located in the **/var/log/messages** file. To prevent the call trace from logging into the kernel log messages, add the **nousb** kernel parameter when the system is booting.

system-config-kdump component, BZ#1077470

In the **Kernel Dump Configuration** window, selecting the **Raw device** option in the **Target settings** tab does not work. To work around this problem, edit the **kdump.conf** file manually.

kernel component, BZ#1087796

An attempt to remove the **bnx2x** module while the **bnx2fc** driver is processing a corrupted frame causes a kernel panic. To work around this problem, shut down any active FCoE interfaces before executing the **modprobe -r bnx2x** command.

kexec-tools component, BZ#1089788

Due to a wrong buffer size calculation in the **makedumpfile** utility, an OOM error could occur with a high probability. As a consequence, the **vmcore** file cannot be captured under certain circumstances. No workaround is currently available.

Chapter 25. Virtualization

libvirt component, BZ#[1095636](#)

SELinux prevents **qemu** from attaching TUN/TAP queues. A multi-queue network interface controller (NIC) is disabled by default in Red Hat Enterprise Linux 7 and should not be turned on.

gnome-boxes component, BZ#[1034354](#)

It is possible to add a Red Hat Enterprise Virtualization Manager 3.2 instance as a source in the **Boxes** tool: the list of virtual machines is displayed, but attempting to connect to a virtual machine can fail with a **Connection failed** message.

spice component, BZ#[1035184](#)

Live migration from Red Hat Enterprise Linux 6.5 to Red Hat Enterprise Linux 7 or from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 7 fails if a SPICE client is connected to the virtual machine.

device-mapper-persistent-data component, BZ#[1057951](#)

The **thin_repair(8)** manual page incorrectly describes that the **thin_repair** utility can repair metadata from an input file. However, this way of using **thin_repair** is not supported and the utility works properly only with metadata from devices. In order to repair metadata from a file, use the **thin_restore** utility instead.

virtio-win component, BZ#[1036341](#)

The current implementation of the QEMU Guest Agent does not allow the Volume Shadow Copy Service (VSS) provider to create shadow copies within a guest. An attempt to create a shadow copy causes an error to be returned. This implementation should only be used as a mean to freeze the file system.

xorg-x11-drv-qxl component, BZ#[1013002](#)

When logging out from a session using **QXL** or **SPICE** on a virtual machine, it is impossible to return to GNOME Display Manager (GDM). Also, GDM cannot be restarted using the **service gdm restart** command.

qemu-kvm component, BZ#[1089610](#)

The minimum VRAM size for a **QXL** device is 16MB in Red Hat Enterprise Linux 6, but it is 4KB in Red Hat Enterprise Linux 7. Consequently, specifying a VRAM size that is less than or equal to 8MB causes the actual VRAM size to differ between the versions, and live migration from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7 fails. To work around this problem, do not specify VRAM sizes that are less than or equal to 8MB on Red Hat Enterprise Linux 6. The default size of 9MB does not cause the described problem.

qemu-kvm component, BZ#[1021483](#)

The serial property for QEMU's **-drive** option is deprecated in Red Hat Enterprise Linux 7. The serial number for virtual disks can still be specified using libvirt or QEMU's **-device** option.

qemu-kvm component, BZ#[983993](#)

Machine types providing the same virtual hardware as Red Hat Enterprise Linux 5 (for example, **rhel-5.4.0**) are no longer available in Red Hat Enterprise Linux 7. Therefore, virtual machines using these machine types cannot be started on or migrated to Red Hat Enterprise Linux 7 without first upgrading them to a newer machine type (for example, **rhel-6.5.0** or **pc**).

xorg-x11-server-utils component, BZ#[1014210](#)

The **xrdb** utility manages the X server resource database. Traditionally, the utility invokes the C preprocessor, which is provided by the *cpp* package, to expand resource files before further operations. However, this feature is disabled by default in Red Hat Enterprise Linux 7 because in most cases, the C preprocessor is not needed. To preserve previous behavior, install *cpp* and execute **xrdb** with the **-cpp** option.

kernel component, BZ#1029295

The **macvtap** driver bypasses the standard linux input stack and diverts the packets sent from user space. This also bypasses the mechanism for delivering packets to the **tcpdump** application. Consequently, the traffic on **macvtap** devices is not captured by default. To work around this problem, execute the following command:

```
tcpdump -i eth0 ether host mac_address
```

With the above command, **tcpdump** filters traffic by mac address assigned to the **macvtap** device. As a result, traffic on **macvtap** devices is captured properly.

virtio-win component, BZ#[1086172](#)

Under certain circumstances, migrating from Red Hat Enterprise Linux 6 to a Red Hat Enterprise Linux 7 host fails with the machine type **rhel6.3.0** or earlier if a **virtio-scsi** device is present. This problem has been observed only with a Window 8 64-bit guest machine. To work around this problem, upgrade the machine type to **rhel6.4.0** or later.

qemu-kvm component, BZ#1071065

The **x86emu** emulator included in the Red Hat Enterprise Linux 5 version of X.Org fails to handle 32-bit prefixes correctly for a number of x86 assembler instructions contained in the vgabios. Consequently, when running a Red Hat Enterprise Linux 5 guest on a Red Hat Enterprise Linux 7 host, the vesa driver does not function on AMD64 and Intel 64 guests, and the X server on the guest machine fails to start. To work around this problem, use the native drivers for cirrus VGA and qxl VGA instead. In order to do this, it might be necessary to edit the **/etc/X11/xorg.conf** file and change the **Driver** entry in the **Device** section from **vesa** to **cirrus** or **qxl**.

qemu-kvm component, BZ#[1043459](#)

When running Red Hat Enterprise Linux 4.9 as a guest virtual machine on a Red Hat Enterprise Linux 7 host, the guest is unable to detect VGA drivers. As a consequence, the GUI does not work. To work around this problem, access the Red Hat Enterprise Linux 4.9 server over an SSH connection using a command line interface, or upgrade the server to Red Hat Enterprise Linux 6 or Red Hat Enterprise Linux 7.

qemu-kvm component, BZ#1071168

The X.Org version included in Red Hat Enterprise Linux 5 does not support automatic configuration of **evdev** devices. Consequently, when running a Red Hat Enterprise Linux 5 guest on a Red Hat Enterprise Linux 7 host, input devices fail to work unless they are configured manually. To work around this problem, edit the **xorg.conf** configuration file on the guest machine according to the following example for the USB tablet emulated by QEMU:

1. Add the following line to the **ServerLayout** section in **xorg.conf**:

```
InputDevice      "usb-tablet" "CorePointer"
```

2. Add the following **InputDevice** section to **xorg.conf**:

```
Section "InputDevice"
    Identifier "usb-tablet"
    Driver "evdev"
    Option "Device" "/dev/input/event2"
EndSection
```

virtio-win component, BZ#[1086084](#)

The **virsh** utility does not support system file-consistent snapshots for Microsoft Windows XP and Windows 2003 virtual guest machines. As a consequence, running the **snapshot-create-as --quiesce --disk-only** command fails. To work around this problem, do not use the **--quiesce** option. Please note that the snapshots will not be consistent with the guest file system.

qemu-kvm component, BZ#[1044979](#)

Microsoft Windows 8.1 and Windows 2012 R2 systems require some CPU features that are not present in all **qemu-kvm** CPU models. Consequently, Microsoft Windows 8.1 and Windows 2012 R2 do not boot if certain CPU models are used, namely Opteron_G1, Conroe, and kvm64. To work around this problem, use CPU models that include the features required by Microsoft Windows 8.1 and Windows 2012 R2.

qemu-kvm component, BZ#731570

The QEMU Enhanced Disk format (QED) for KVM guest virtual machines is not supported in Red Hat Enterprise Linux 7. Use the qcow2 image format instead.

seabios component, BZ#[1034072](#)

The **SeaBIOS** application runs in real mode for compatibility with BIOS interfaces. This limits the amount of memory available. As a consequence, **SeaBIOS** is only able to handle a limited number of disks. Currently, the supported number of disks is:

- ▶ **virtio-scsi** — 64
- ▶ **virtio-blk** — 4
- ▶ **ahci/sata** — 24 (4 controllers with all 6 ports connected)
- ▶ **usb-storage** — 4

As a workaround for this problem, when attaching a large number of disks to your virtual machine, make sure that your system disk has a small **pci** slot number, so **SeaBIOS** sees it first when scanning the **pci** bus. It is also recommended to use the **virtio-scsi** device instead of **virtio-blk** as the per-disk memory overhead is smaller.

libvirt component, BZ#1100588

When installing Red Hat Enterprise Linux 7 from other sources than network, the network devices are not specified by default in the interface configuration files. As a consequence, creating a bridge by using the **iface-bridge** command in the **virsh** utility fails with an error message. To work around the problem, add the **DEVICE=** lines in the **/etc/sysconfig/network-scripts/ifcfg-*** files.

kernel component, BZ#1075857

The emulated Small Computer System Interface (SCSI) driver, **sym53c895**, is not supported in Red Hat Enterprise Linux 7. As a consequence, a Red Hat Enterprise Linux 7 guest running under the Xen virtual machine monitor is currently unable to find devices, volumes, and images attached to the Red Hat Enterprise Linux 7 guest through emulated SCSI drivers. To work around this problem, it is possible to use paravirtualized drivers, **xen-blkfront**, as they are supported in Red Hat Enterprise Linux 7 and bring significant performance improvements. To do this, edit the Xen domain configuration file, usually found in the **/etc/xen/** directory, and specify the disk in question as for example **file:/root/raw.img,xvdc,w** instead of **file:/root/raw.img,sdc,w**.

grub2 component, BZ#1045127

When using the serial console through KVM, holding down an arrow key for an extended period of time to navigate in the GRUB 2 menu results in erratic behavior. To work around this problem, avoid the rapid input caused by holding down an arrow key for a longer time.

Chapter 26. Deployment and Tools

systemd component, BZ#978955

When attempting to start, stop or restart a service or unit using the **systemctl** **[start|stop|restart] NAME** command, no message is displayed to inform the user whether the action has been successful.

systemd component, BZ#968401

The **/etc/rc.d/rc.local** file does not have executable permissions in Red Hat Enterprise Linux 7. If commands are added to the **/etc/rc.d/rc.local** file, the file has to be made executable afterwards.

By default, **/etc/rc.d/rc.local** does not have executable permissions because if these permissions are detected, the system has to wait until the network is up before the bootup can be finished.

flightrecorder component, BZ#1049701

The *flightrecorder* package is currently not included in Red Hat Enterprise Linux 7.

Chapter 27. Clustering

resource-agents component, BZ#[1077888](#)

The **CTDB** agent used to implement an High Availability **samba** does not work as expected in Red Hat Enterprise Linux 7. If you wish to configure clustered Samba for Red Hat Enterprise Linux 7, follow the steps in this Knowledge Base article:

<https://access.redhat.com/site/articles/912273>

Chapter 28. Networking

iptables component, BZ#1018135

Red Hat Enterprise Linux 7 introduces the *iptables* packages, which replace the *iptables_jf* packages shipped in Red Hat Enterprise Linux 6. All users of *iptables* are advised to update their scripts because the syntax of this version differs from *iptables_jf*.

rsync component, BZ#1082496

The **rsync** utility cannot be run as a socket-activated service because the **rsyncd@.service** file is missing from the *rsync* package. Consequently, the **systemctl start systemd.socket** command does not work. However, running **rsync** as a daemon by executing the **systemctl start systemd.service** command works as expected.

openssl component, BZ#1062656

It is not possible to connect to any Wi-Fi Protected Access (WPA) Enterprise Access Point (AP) that requires MD5-signed certificates. To work around this problem, copy the **wpa_supplicant.service** file from the **/usr/lib/systemd/system/** directory to the **/etc/systemd/system/** directory and add the following line to the **Service** section of the file:

```
Environment="OPENSSL_ENABLE_MD5_VERIFY"
```

Then run the **systemctl daemon-reload** command as root to reload the service file.



Important

Note that MD5 certificates are highly insecure and Red Hat does not recommend using them.

bind component, BZ#1004300

Previously, **named-chroot.service** was setting up the **chroot** environment for the **named** daemon by mounting necessary files and directories to the **/var/named/chroot/** path before starting the daemon. However, if the start up of the daemon failed, the mounts remained mounted. As a consequence, the **chroot** environment was corrupted. This also affected **named-sdb-chroot.service**, which was using the same **chroot** path. With this update, **named-chroot.service** and **named-sdb-chroot.service** have been modified and the **chroot** set up code has been separated into two new **systemd** services, **named-chroot-setup.service** and **named-sdb-chroot-setup.service**. In addition, the **named-sdb** daemon now uses its own **chroot** path, **/var/named/chroot_sdb/**. Also, **named-sdb** daemon has been removed from the *bind-chroot* package and is now included in its own *bind-sdb-chroot* subpackage. Users who use **named-sdb** in the **chroot** environment are advised to install the *bind-sdb-chroot* package.

bind-dyndb-ldap component, BZ#1078295

The **bind-dyndb-ldap** plug-in does not fully support the DNS64 server. As a consequence, the **BIND** daemon configured with DNS64 terminates unexpectedly when a DNS64 query is processed by **bind-dyndb-ldap**. To work around this problem, disable DNS64 in the **named.conf** file. The whole section concerning DNS64 can be commented out.

openswitch component, BZ#1066493

In certain cases, when connecting two network interface controllers (NIC) that use the **ixgbe** driver, the TCP stream throughput does not exceed 8.4 GB. This problem manifests itself both on a NIC to NIC level, although to a very limited degree, as well as in combination with virtual machines running on top of an **openvswitch** bridge.

vsftpd component, [BZ#1058712](#)

The vsftpd daemon does not currently support ciphers suites based on the ECDHE key-assignment protocol. Consequently, when vsftpd is configured to use such suite, the connection is refused with **no shared cipher** SSL alert.

fcoe-utils component, [BZ#1049200](#)

The **-m vn2vn** option of the **fcoeadm** command does not work correctly, and Fabric mode is always used instead of vn2vn. As a consequence, a vn2vn instance cannot be created using **fcoeadm**, and the port state is offline instead of online. To work around this problem, modify the **sysfs** file manually to create a vn2vn link.

NetworkManager component, BZ#1030947

The **brctl addbr name** command, which is used for creating a new instance of the Ethernet bridge, also brings the interface up. Consequently, the **brctl delbr name** command does not delete the instance of the Ethernet bridge because the network interface corresponding to the bridge is not down. To work around the problem:

- Either bring the instance down by using the **ip link set dev name down** command before running the **brctl delbr name** command,
- Or use the **ip link del name** command for deleting the instance.

Chapter 29. Authentication and Interoperability

sssd component, BZ#1081046

The **accountExpires** attribute **SSSD** uses to see whether an account has expired is not replicated to the Global Catalog by default. As a result, users with expired accounts can be allowed to log in when using GSSAPI authentication. To work around this problem, the Global Catalog support can be disabled by specifying **ad_enable_gc=False** in the **sssd.conf** file. With this setting, users with expired accounts will be denied access when using GSSAPI authentication. Note that **SSSD** connects to each LDAP server individually in this scenario, which can increase the connection count.

ipa component, BZ#1004156

When DNS support is being added for an Identity Management server (for example, by using the **ipa-dns-install** or by using the **--setup-dns** flag in **ipa-server-install** or **ipa-replica-install**), the script adds a hostname of a new Identity Management DNS server to the list of name servers in the primary Identity Management DNS zone (via DNS NS record). However, it does not add the DNS name server record to other DNS zones served by the Identity Management. As a consequence, the list of name servers in the non-primary DNS zones has only a limited set of Identity Management name servers serving the DNS zone (only one, without user intervention). When the limited set of Identity Management name servers is not available, these DNS zones are not resolvable. To work around this problem, manually add new DNS name server records to all non-primary DNS zones when a new Identity Management replica is being added. Also, manually remove such DNS name server records when the replica is being decommissioned. Non-primary DNS zones can maintain higher availability by having a manually maintained set of Identity Management name servers serving it.

ipa component, BZ#971384

The default **Unlock user accounts** permission does not include the **nsaccountlock** attribute, which is necessary for a successful unlocking of a user entry. Consequently, the privileged user with this permission assigned cannot unlock another user, and errors like the following are displayed:

```
ipa: ERROR: Insufficient access: Insufficient 'write' privilege to the
'nsAccountLock' attribute of entry
'uid=user,cn=users,cn=accounts,dc=example,dc=com'.
```

To work around this problem, add **nsaccountlock** to the list of allowed attributes in the aforementioned permission by running the following command:

```
~]# ipa permission-mod "Unlock user accounts" --attrs=
{krbLastAdminUnlock,krbLoginFailedCount,nsaccountlock}
```

As a result, users with the **Unlock user accounts** permission assigned can unlock other users.

ipa component, BZ#970618

The Identity Management Kerberos driver does not actively update the default PAC types for Kerberos tickets issued by Identity Management Kerberos key distribution center (KDC). Consequently, if the default list of PAC types is changed in the Identity Management configuration, the Identity Management Kerberos KDC does not generate the configured PAC types for the issued tickets until the KDC is restarted. Changing the default list can be, for example, performed by running the following command:

```
~]# ipa config-mod --pac-type NEW-PAC-TYPES
```

To work around this problem, restart the Identity Management Kerberos KDC service on all Identity Management servers. As a result, Identity Management Kerberos KDC generates the configured PAC types as set in the Identity Management server configuration.

ipa component, BZ#[973195](#)

There are multiple problems across different tools used in the Identity Management installation, which prevents installation of CA-less certificates with intermediate certificate authority (CA). One of the errors is that incorrect trust flags are assigned to the intermediate CA certificate when importing a PKCS#12 file. Consequently, the Identity Management server installer fails due to an incomplete trust chain that is returned for Identity Management services. There is no known workaround, CA-less certificates must not contain intermediate CA in their trust chain.

ipa component , BZ#[988473](#)

Access control to lightweight directory access protocol (LDAP) objects representing trust with Active Directory (AD) is given to the **Trusted Admins** group in Identity Management. In order to establish the trust, the Identity Management administrator should belong to a group which is a member of the "Trusted Admins" group and this group should have relative identifier (RID) 512 assigned. To ensure this, run the **ipa-adtrust-install** command and then the **ipa group-show admins --all** command to verify that the "ipantsecurityidentifier" field contains a value ending with the "-512" string. If the field does not end with "-512", use the **ipa group-mod admins --setattr=ipantsecurityidentifier=SID** command, where SID is the value of the field from the **ipa group-show admins --all** command output with the last component value (-XXXX) replaced by the "-512" string.

ipa component, BZ#[1084018](#)

Red Hat Enterprise Linux 7 contains an updated version of **slapi-nis**, a Directory Server plugin, which allows users of Identity Management and the Active Directory service to authenticate on legacy clients. However, the **slapi-nis** component only enables identity and authentication services, but does not allow users to change their password. As a consequence, users logged to legacy clients via **slapi-nis** compatibility tree can change their password only via the Identity Management Server Self-Service Web UI page or directly in Active Directory.

ipa component, BZ#[1060349](#)

The **ipa host-add** command does not verify the existence of AAAA records. As a consequence, **ipa host-add** fails if no A record is available for the host, although an AAAA record exists. To work around this problem, run **ipa host-add** with the **--force** option.

ipa component, BZ#[1081626](#)

An IPA master is uninstalled while SSL certificates for services other than IPA servers are tracked by the **certmonger** service. Consequently, an unexpected error can occur, and the uninstallation fails. To work around this problem, start **certmonger**, and run the **ipa-getcert** command to list the tracked certificates. Then run the **ipa-getcert stop-tracking -i <Request ID>** command to stop **certmonger** from tracking the certificates, and run the IPA uninstall script again.

ipa component, BZ#[1088683](#)

The **ipa-client-install** command does not process the **--preserve-sssd** option

correctly when generating the IPA domain configuration in the **sssd.conf** file. As a consequence, the original configuration of the IPA domain is overwritten. To work around this problem, review **sssd.conf** after running **ipa-client-install** to identify and manually fix any unwanted changes.

certmonger component, BZ#[996581](#)

The directory containing a private key or certificate can have an incorrect SELinux context. As a consequence, the **ipa-getcert request -k** command fails, and an unhelpful error message is displayed. To work around this problem, set the SELinux context on the directory containing the certificate and the key to **cert_t**. You can resubmit an existing certificate request by running the **ipa-getcert resubmit -i <Request ID>** command.

sssd component, BZ#[1103249](#)

Under certain circumstances, the algorithm in the Privilege Attribute Certificate (PAC) responder component of the System Security Services Daemon (SSSD) does not effectively handle users who are members of a large number of groups. As a consequence, logging from Windows clients to Red Hat Enterprise Linux clients with Kerberos single sign-on (SSO) can be noticeably slow. There is currently no known workaround available.

ipa component, BZ#[1033357](#)

The **ipactl restart** command requires the directory server to be running. Consequently, if this condition is not met, **ipactl restart** fails with an error message. To work around this problem, use the **ipactl start** command to start the directory server before executing **ipactl restart**. Note that the **ipactl status** command can be used to verify if the directory server is running.

pki-core component, BZ#1085105

The certificate subsystem fails to install if the system language is set to Turkish. To work around this problem, set the system language to English by putting the following line in the **/etc/sysconfig/i18n** file:

```
LANG="en_US.UTF-8"
```

Also, remove any other "LANG=" entries in **/etc/sysconfig/i18n**, then reboot the system. After reboot, you can successfully run **ipa-server-install**, and the original contents of **/etc/sysconfig/i18n** may be restored.

ipa component, BZ#[1020563](#)

The **ipa-server-install** and **ipa-replica-install** commands replace the list of NTP servers in the **/etc/ntp.conf** file with Red Hat Enterprise Linux default servers. As a consequence, NTP servers configured before installing IPA are not contacted, and servers from **rhel.pool.ntp.org** are contacted instead. If those default servers are unreachable, the IPA server does not synchronize its time via NTP. To work around this problem, add any custom NTP servers to **/etc/ntp.conf**, and remove the default Red Hat Enterprise Linux servers if required. The configured servers are now used for time synchronization after restarting the NTP service by running the **systemctl restart ntpd.service** command.

ipa component, BZ#[1093159](#)

The "Modify Sudo rule" permission does not include the **memberhost** and **externalhost** attributes. As a consequence, when a user is logged in with the "Modify Sudo Rule" permission, commands such as the following:

```
~]# ipa sudorule-add-host sudorule --hosts=ipa-replica.ipa.example
```

and equivalent web user interface actions fail with the following error message:

```
Insufficient access: Insufficient 'write' privilege to the 'externalHost'
attribute of entry
'ipauniqueid=$UUID,cn=sudorules,cn=sudo,dc=ipa,dc=example'.
```

As a workaround, add the **memberhost** and **externalhost** attributes to the "Modify Sudo rule" permission, either using the web user interface, or by running the following command:

```
~]# ipa permission-mod "Modify Sudo rule" --attrs=
{description,ipaenabledflag,usercategory,hostcategory,cmdcategory,ipasud
orunasusercategory,ipasudorunasgroupcategory,externaluser,ipasudorunasext
user,
ipasudorunasextgroup,memberdenycommand,memberallowcommand,memberuser,memberhost,e
xternalhost}
```

After running the command, users logged in with the **Modify Sudo rule** permission can add member hosts to **sudo** rules.

gnutls component, BZ#[1084080](#)

The **gnutls** utility fails to generate a non-encrypted private key when the user enters an empty password. To work around this problem, use the **certtool** command with the **password** option as follows:

```
~]$ certtool --generate-privkey --pkcs8 --password "" --outfile pkcs8.key
```

Chapter 30. Security

polycoreutils component, BZ#[1082676](#)

Due to a bug in the **fixfiles** scripts, if the **exclude_dirs** file is defined to exclude directories from relabeling, running the **fixfiles restore** command applies incorrect labels on numerous files on the system.

SELinux component

Note that a number of daemons that were previously not confined are confined in Red Hat Enterprise Linux 7.

Chapter 31. Entitlement

subscription-manager component, BZ#910345

The default **firstboot** behavior is to prompt for Subscription Manager or Subscription Asset Manager (SAM) details. It no longer offers a path to use the Red Hat Network Classic registration tool.

rhn-client-tools component, BZ#910345

The **rhn-client-tools** utility is no longer configured by default to communicate with `xmlrpc.rhn.redhat.com`, instead, it prompts for the user's Red Hat Satellite or Red Hat Proxy details to be entered.

Chapter 32. Desktop

spice component, BZ#[1030024](#)

Video playback on a Red Hat Enterprise Linux 7 guest with GNOME Shell is sometimes not detected as a video stream by **spice-server**. The video stream is therefore not compressed in such a case.

kde component

A companion GUI to unixODBC based on KDE is no longer maintained by upstream. There is an independent project, [unixODBC-GUI-Qt](#), however, no stable release is available at the moment. Therefore, no supported GUI for unixODBC is available for Red Hat Enterprise Linux 7.

mutter component, BZ#[861507](#)

Support for quad-buffered OpenGL stereo visual effects in compositing window managers, such as **Mutter**, is missing in Red Hat Enterprise Linux 7. NVIDIA drivers currently do not support stereo visual effects when a compositing window manager is running.

xorg-x11-drv-nouveau component, BZ#915857

Under some circumstances, reconnecting a monitor from one GPU port to another while the system is running can render the X server completely unresponsive to mouse and keyboard inputs. Sometimes, the monitor does not turn on in the described scenario. Also, the same problem occurs in runlevel 3, that is without the X Window system running.

gobject-introspection component, BZ#[1076414](#)

The **gobject-introspection** library is not available in a 32-bit multilib package. Users who wish to compile 32-bit applications that rely on GObject introspection or libraries that use it, such as **GTK+** or **GLib**, should use the *mock* package to set up a build environment for their applications.

Revision History

| | | |
|-----------------------|------------------------|-------------------------|
| Revision 0.0-3 | Tue Jun 10 2014 | Eliška Slobodová |
|-----------------------|------------------------|-------------------------|

Release of the Red Hat Enterprise Linux 7.0 Release Notes.

| | | |
|-----------------------|------------------------|-------------------------|
| Revision 0.0-1 | Thu Dec 11 2013 | Eliška Slobodová |
|-----------------------|------------------------|-------------------------|

Release of the Red Hat Enterprise Linux 7.0 Beta Release Notes.