



Red Hat Enterprise Linux 8.3

8.3 Release Notes

Release Notes for Red Hat Enterprise Linux 8.3

Red Hat Enterprise Linux 8.3 8.3 Release Notes

Release Notes for Red Hat Enterprise Linux 8.3

Legal Notice

Copyright © 2023 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 8.3 and document known problems in this release, as well as notable bug fixes, Technology Previews, deprecated functionality, and other details.

Table of Contents

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	5
CHAPTER 1. OVERVIEW	6
Installer and image creation	6
RHEL for Edge	6
Infrastructure services	6
Security	6
Dynamic programming languages, web and database servers	6
Compiler toolsets	7
Identity Management	7
The web console	7
Virtualization	7
Desktop and graphics	7
In-place upgrade and OS conversion	7
OpenJDK 11 is now available	8
Additional resources	8
Red Hat Customer Portal Labs	9
CHAPTER 2. ARCHITECTURES	10
CHAPTER 3. DISTRIBUTION OF CONTENT IN RHEL 8	11
3.1. INSTALLATION	11
3.2. REPOSITORIES	11
3.3. APPLICATION STREAMS	12
3.4. PACKAGE MANAGEMENT WITH YUM/DNF	12
CHAPTER 4. RHEL 8.3.1 RELEASE	13
4.1. NEW FEATURES	13
CHAPTER 5. RHEL 8.3.0 RELEASE	17
5.1. NEW FEATURES	17
5.1.1. Installer and image creation	17
5.1.2. RHEL for Edge	19
5.1.3. Software management	19
5.1.4. Shells and command-line tools	20
5.1.5. Infrastructure services	21
5.1.6. Security	24
5.1.7. Networking	32
5.1.8. Kernel	34
5.1.9. File systems and storage	38
5.1.10. High availability and clusters	40
5.1.11. Dynamic programming languages, web and database servers	42
5.1.12. Compilers and development tools	47
5.1.13. Identity Management	53
5.1.14. Desktop	60
5.1.15. Graphics infrastructures	61
5.1.16. The web console	62
5.1.17. Red Hat Enterprise Linux System Roles	62
5.1.18. Virtualization	64
5.1.19. RHEL in cloud environments	70
5.1.20. Containers	70
5.2. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS	71

New kernel parameters	71
Updated kernel parameters	73
New /proc/sys/fs parameters	75
5.3. DEVICE DRIVERS	76
5.3.1. New drivers	76
Network drivers	76
Graphics drivers and miscellaneous drivers	76
5.3.2. Updated drivers	77
Network driver updates	77
Storage driver updates	78
Graphics and miscellaneous driver updates	78
5.4. BUG FIXES	78
5.4.1. Installer and image creation	78
5.4.2. Software management	80
5.4.3. Shells and command-line tools	80
5.4.4. Security	80
5.4.5. Networking	83
5.4.6. Kernel	83
5.4.7. High availability and clusters	85
5.4.8. Dynamic programming languages, web and database servers	86
5.4.9. Compilers and development tools	86
5.4.10. Identity Management	89
5.4.11. Graphics infrastructures	91
5.4.12. Virtualization	91
5.4.13. Containers	92
5.5. TECHNOLOGY PREVIEWS	92
5.5.1. Networking	92
5.5.2. Kernel	94
5.5.3. File systems and storage	95
5.5.4. High availability and clusters	97
5.5.5. Identity Management	98
5.5.6. Desktop	99
5.5.7. Graphics infrastructures	99
5.5.8. Red Hat Enterprise Linux System Roles	100
5.5.9. Virtualization	100
5.5.10. Containers	101
5.6. DEPRECATED FUNCTIONALITY	102
5.6.1. Installer and image creation	102
5.6.2. Software management	103
5.6.3. Shells and command-line tools	103
5.6.4. Infrastructure services	104
5.6.5. Security	104
5.6.6. Networking	105
5.6.7. Kernel	105
5.6.8. File systems and storage	106
5.6.9. Identity Management	107
5.6.10. Desktop	108
5.6.11. Graphics infrastructures	109
5.6.12. The web console	109
5.6.13. Red Hat Enterprise Linux System Roles	109
5.6.14. Virtualization	109
5.6.15. Containers	110
5.6.16. Deprecated packages	110

5.7. KNOWN ISSUES	110
5.7.1. Installer and image creation	111
5.7.2. Subscription management	114
5.7.3. Infrastructure services	114
5.7.4. Security	115
5.7.5. Networking	120
5.7.6. Kernel	121
5.7.7. File systems and storage	126
5.7.8. Dynamic programming languages, web and database servers	127
5.7.9. Identity Management	127
5.7.10. Desktop	129
5.7.11. Graphics infrastructures	130
5.7.12. The web console	131
5.7.13. Red Hat Enterprise Linux System Roles	132
5.7.14. Virtualization	132
5.7.15. RHEL in cloud environments	133
5.7.16. Supportability	134
5.7.17. Containers	135
CHAPTER 6. INTERNATIONALIZATION	136
6.1. RED HAT ENTERPRISE LINUX 8 INTERNATIONAL LANGUAGES	136
6.2. NOTABLE CHANGES TO INTERNATIONALIZATION IN RHEL 8	136
APPENDIX A. LIST OF TICKETS BY COMPONENT	138
APPENDIX B. REVISION HISTORY	146

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Please let us know how we could make it better. To do so:

- For simple comments on specific passages, make sure you are viewing the documentation in the Multi-page HTML format. Highlight the part of text that you want to comment on. Then, click the **Add Feedback** pop-up that appears below the highlighted text, and follow the displayed instructions.
- For submitting more complex feedback, create a Bugzilla ticket:
 1. Go to the [Bugzilla](#) website.
 2. As the Component, use **Documentation**.
 3. Fill in the **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation.
 4. Click **Submit Bug**.

CHAPTER 1. OVERVIEW

Installer and image creation

In RHEL 8.3, you can configure a root password and create a user account before you begin the installation. Previously, you configured a root password and created a user account after you began the installation process. You can also create customized images based on a much more reliable backend and also push images to clouds through the RHEL web console.

RHEL for Edge

RHEL 8.3 introduces **RHEL for Edge** for remotely installing RHEL on Edge servers. RHEL for Edge is an rpm-ostree image that you can compose using Image Builder. You can install the image using a Kickstart file and then manage the image to include image updates and to roll back an image to a previous functional state.

Following are RHEL for Edge key highlights:

- Atomic upgrades, where the state of each update is known and no changes are seen until you reboot the device.
- Custom health checks and intelligent rollbacks to ensure resiliency.
- Container-focused workflows, where you can separate core OS updates from the application updates, and test and deploy different versions of applications.
- Optimized OTA payloads for low-bandwidth environments.

For more information, see [Section 5.1.2, “RHEL for Edge”](#).

Infrastructure services

The **Tuned** system tuning tool has been rebased to version 2.13, which adds support for architecture-dependent tuning and multiple include directives.

Security

RHEL 8.3 provides Ansible roles for automated deployments of Policy-Based Decryption (PBD) solutions using **Clevis** and **Tang**, and this version of the **rhel-system-roles** package also contains an Ansible role for RHEL logging through **Rsyslog**.

The **scap-security-guide** packages have been rebased to version 0.1.50, and **OpenSCAP** has been rebased to version 1.3.3. These updates provide substantial improvements, including a profile aligned with the CIS RHEL 7 Benchmark v2.2.0 and a profile aligned with the Health Insurance Portability and Accountability Act (HIPAA) that is required by North-American healthcare organizations.

With this update, you can now generate result-based remediation roles from tailored profiles using the **SCAP Workbench** tool.

The **USBGuard** framework now provides its own SELinux policy, it notifies desktop users in GUI, and the version 0.7.8 contains many other improvements and bug fixes.

Dynamic programming languages, web and database servers

Later versions of the following components are now available as new module streams:

- **nginx 1.18**
- **Node.js 14**
- **Perl 5.30**

- **PHP 7.4**
- **Ruby 2.7**

The following components have been updated in RHEL 8.3:

- **Git** to version 2.27
- **Squid** to version 4.11

See [Section 5.1.11, “Dynamic programming languages, web and database servers”](#) for more information.

Compiler toolsets

The following compiler toolsets have been updated in RHEL 8.3:

- **GCC Toolset 10**
- **LLVM Toolset 10.0.1**
- **Rust Toolset 1.45.2**
- **Go Toolset 1.14.7**

See [Section 5.1.12, “Compilers and development tools”](#) for more information.

Identity Management

The Rivest Cipher 4 (RC4) cipher suite, the default encryption type for users, services, and trusts between Active Directory (AD) domains in an AD forest, has been deprecated in RHEL 8. For compatibility reasons, this update introduces a new cryptographic subpolicy **AD-SUPPORT** to enable support for the deprecated RC4 encryption type. The new subpolicy allows you to use RC4 with RHEL Identity Management (IdM) and SSSD Active Directory integration solutions.

See [Section 5.1.13, “Identity Management”](#) for more information.

The web console

The web console provides an option to switch between administrative access and limited access from inside of a user session.

Virtualization

Virtual machines (VMs) hosted on IBM Z hardware can now use the IBM Secure Execution feature. This makes the VMs resistant to attacks if the host is compromised, and also prevents untrusted hosts from obtaining information from the VM. In addition, DASD devices can now be assigned to VMs on IBM Z.

Desktop and graphics

You can now use the GNOME desktop on IBM Z systems.

The **Direct Rendering Manager** (DRM) kernel graphics subsystem has been rebased to upstream Linux kernel version 5.6. This version provides a number of enhancements over the previous version, including support for new GPUs and APUs, and various driver updates.

See [Section 5.1.14, “Desktop”](#) and [Section 5.1.15, “Graphics infrastructures”](#) for further details.

In-place upgrade and OS conversion

In-place upgrade from RHEL 7 to RHEL 8

The supported in-place upgrade paths currently are:

- From RHEL 7.8 to RHEL 8.2 on the 64-bit Intel, IBM POWER 8 (little endian), and IBM Z architectures
- From RHEL 7.6 to RHEL 8.2 on architectures that require kernel version 4.14: IBM POWER 9 (little endian) and IBM Z (Structure A)
- From RHEL 7.7 to RHEL 8.2 on systems with SAP HANA.

To ensure your system remains supported after upgrading to RHEL 8.2, either update to the latest RHEL 8.3 version or enable the RHEL 8.2 Extended Update Support (EUS) repositories. On systems with SAP HANA, enable the RHEL 8.2 Update Services for SAP Solutions (E4S) repositories.

For more information, see [Supported in-place upgrade paths for Red Hat Enterprise Linux](#) . For instructions on performing an in-place upgrade, see [Upgrading from RHEL 7 to RHEL 8](#) .

Notable enhancements include:

- **Leapp** now supports user input by generating true/false questions to determine how to proceed with the upgrade.
- You can now upgrade multiple hosts simultaneously using the Satellite web UI.
- The in-place upgrade is now supported for on-demand instances on AWS and Microsoft Azure, using Red Hat Update Infrastructure (RHUI).
- With the release of the [RHBA-2021:0569](#) advisory, you can create custom scripts for the **Leapp** pre-upgrade report. See [Automating your Red Hat Enterprise Linux pre-upgrade report workflow](#) for details.

In-place upgrade from RHEL 6 to RHEL 8

To upgrade from RHEL 6.10 to RHEL 8.2, follow instructions in [Upgrading from RHEL 6 to RHEL 8](#) .

Conversion from a different Linux distribution to RHEL

If you are using CentOS Linux 8 or Oracle Linux 8, you can convert your operating system to RHEL 8 using the Red Hat-supported **Convert2RHEL** utility. For more information, see [Converting from an RPM-based Linux distribution to RHEL](#).

If you are using an earlier version of CentOS Linux or Oracle Linux, namely versions 6 or 7, you can convert your operating system to RHEL and then perform an in-place upgrade to RHEL 8. Note that CentOS Linux 6 and Oracle Linux 6 conversions use the unsupported **Convert2RHEL** utility. For more information on unsupported conversions, see [How to convert from CentOS Linux 6 or Oracle Linux 6 to RHEL 6](#).

For information regarding how Red Hat supports conversions from other Linux distributions to RHEL, see the [Convert2RHEL Support Policy document](#) .

OpenJDK 11 is now available

New version of Open Java Development Kit (OpenJDK) is now available. For more information about the features introduced in this release and changes in the existing functionality, see [OpenJDK 11 documentation](#).

Additional resources

- **Capabilities and limits** of Red Hat Enterprise Linux 8 as compared to other versions of the system are available in the Knowledgebase article [Red Hat Enterprise Linux technology capabilities and limits](#).

supplements and links:

- Information regarding the Red Hat Enterprise Linux **life cycle** is provided in the [Red Hat Enterprise Linux Life Cycle](#) document.
- The [Package manifest](#) document provides a **package listing** for RHEL 8.
- Major **differences between RHEL 7 and RHEL 8** are documented in [Considerations in adopting RHEL 8](#).
- Instructions on how to perform an **in-place upgrade from RHEL 7 to RHEL 8** are provided by the document [Upgrading from RHEL 7 to RHEL 8](#).
- The **Red Hat Insights** service, which enables you to proactively identify, examine, and resolve known technical issues, is now available with all RHEL subscriptions. For instructions on how to install the Red Hat Insights client and register your system to the service, see the [Red Hat Insights Get Started](#) page.

Red Hat Customer Portal Labs

Red Hat Customer Portal Labs is a set of tools in a section of the Customer Portal available at <https://access.redhat.com/labs/>. The applications in Red Hat Customer Portal Labs can help you improve performance, quickly troubleshoot issues, identify security problems, and quickly deploy and configure complex applications. Some of the most popular applications are:

- [Registration Assistant](#)
- [Product Life Cycle Checker](#)
- [Kickstart Generator](#)
- [Kickstart Converter](#)
- [Red Hat Enterprise Linux Upgrade Helper](#)
- [Red Hat Satellite Upgrade Helper](#)
- [Red Hat Code Browser](#)
- [JVM Options Configuration Tool](#)
- [Red Hat CVE Checker](#)
- [Red Hat Product Certificates](#)
- [Load Balancer Configuration Tool](#)
- [Yum Repository Configuration Helper](#)

CHAPTER 2. ARCHITECTURES

Red Hat Enterprise Linux 8.3 is distributed with the kernel version 4.18.0-240, which provides support for the following architectures:

- AMD and Intel 64-bit architectures
- The 64-bit ARM architecture
- IBM Power Systems, Little Endian
- 64-bit IBM Z

Make sure you purchase the appropriate subscription for each architecture. For more information, see [Get Started with Red Hat Enterprise Linux - additional architectures](#) . For a list of available subscriptions, see [Subscription Utilization](#) on the Customer Portal.

CHAPTER 3. DISTRIBUTION OF CONTENT IN RHEL 8

3.1. INSTALLATION

Red Hat Enterprise Linux 8 is installed using ISO images. Two types of ISO image are available for the AMD64, Intel 64-bit, 64-bit ARM, IBM Power Systems, and IBM Z architectures:

- Binary DVD ISO: A full installation image that contains the BaseOS and AppStream repositories and allows you to complete the installation without additional repositories.



NOTE

The Binary DVD ISO image is larger than 4.7 GB, and as a result, it might not fit on a single-layer DVD. A dual-layer DVD or USB key is recommended when using the Binary DVD ISO image to create bootable installation media. You can also use the Image Builder tool to create customized RHEL images. For more information about Image Builder, see the [Composing a customized RHEL system image](#) document.

- Boot ISO: A minimal boot ISO image that is used to boot into the installation program. This option requires access to the BaseOS and AppStream repositories to install software packages. The repositories are part of the Binary DVD ISO image.

See the [Performing a standard RHEL 8 installation](#) document for instructions on downloading ISO images, creating installation media, and completing a RHEL installation. For automated Kickstart installations and other advanced topics, see the [Performing an advanced RHEL 8 installation](#) document.

3.2. REPOSITORIES

Red Hat Enterprise Linux 8 is distributed through two main repositories:

- BaseOS
- AppStream

Both repositories are required for a basic RHEL installation, and are available with all RHEL subscriptions.

Content in the BaseOS repository is intended to provide the core set of the underlying OS functionality that provides the foundation for all installations. This content is available in the RPM format and is subject to support terms similar to those in previous releases of RHEL. For a list of packages distributed through BaseOS, see the [Package manifest](#).

Content in the Application Stream repository includes additional user space applications, runtime languages, and databases in support of the varied workloads and use cases. Application Streams are available in the familiar RPM format, as an extension to the RPM format called *modules*, or as Software Collections. For a list of packages available in AppStream, see the [Package manifest](#).

In addition, the CodeReady Linux Builder repository is available with all RHEL subscriptions. It provides additional packages for use by developers. Packages included in the CodeReady Linux Builder repository are unsupported.

For more information about RHEL 8 repositories, see the [Package manifest](#).

3.3. APPLICATION STREAMS

Red Hat Enterprise Linux 8 introduces the concept of Application Streams. Multiple versions of user space components are now delivered and updated more frequently than the core operating system packages. This provides greater flexibility to customize Red Hat Enterprise Linux without impacting the underlying stability of the platform or specific deployments.

Components made available as Application Streams can be packaged as modules or RPM packages and are delivered through the AppStream repository in RHEL 8. Each Application Stream component has a given life cycle, either the same as RHEL 8 or shorter. For details, see [Red Hat Enterprise Linux Life Cycle](#).

Modules are collections of packages representing a logical unit: an application, a language stack, a database, or a set of tools. These packages are built, tested, and released together.

Module streams represent versions of the Application Stream components. For example, several streams (versions) of the PostgreSQL database server are available in the **postgresql** module with the default **postgresql:10** stream. Only one module stream can be installed on the system. Different versions can be used in separate containers.

Detailed module commands are described in the [Installing, managing, and removing user-space components](#) document. For a list of modules available in AppStream, see the [Package manifest](#).

3.4. PACKAGE MANAGEMENT WITH YUM/DNF

On Red Hat Enterprise Linux 8, installing software is ensured by the **YUM** tool, which is based on the **DNF** technology. We deliberately adhere to usage of the **yum** term for consistency with previous major versions of RHEL. However, if you type **dnf** instead of **yum**, the command works as expected because **yum** is an alias to **dnf** for compatibility.

For more details, see the following documentation:

- [Installing, managing, and removing user-space components](#)
- [Considerations in adopting RHEL 8](#)

CHAPTER 4. RHEL 8.3.1 RELEASE

Red Hat makes Red Hat Enterprise Linux 8 content available quarterly, in between minor releases (8.Y). The quarterly releases are numbered using the third digit (8.Y.1). The new features in the RHEL 8.3.1 release are described below.

4.1. NEW FEATURES

Flatpak packages for several desktop applications

Flatpak is a system for running graphical applications as containers. Using Flatpak, you can install and update an application independently of the host operating system.

This update provides Flatpak container images of the following applications in the Red Hat Container Catalog:

Application name	Flatpak container ID
Firefox	org.mozilla.firefox
GIMP	org.gimp.GIMP
Inkscape	org.inkscape.Inkscape
Thunderbird	org.mozilla.Thunderbird

To install Flatpak containers available in the Red Hat Container Catalog, use the following procedure:

1. Make sure that the latest version of the Flatpak client is installed on your system:

```
# yum update flatpak
```

2. Enable the RHEL Flatpak repository:

```
# flatpak remote-add rhel https://flatpaks.redhat.io/rhel.flatpakrepo
```

3. Provide the credentials for your RHEL account:

```
# podman login registry.redhat.io
```

By default, Podman saves the credentials only until the user logs out.

4. Optional: Save your credentials permanently:

```
$ cp $XDG_RUNTIME_DIR/containers/auth.json \
  $HOME/.config/flatpak/oci-auth.json
```

5. Install the Flatpak container image:

```
$ flatpak install rhel container-id
```

(JIRA:RHELPLAN-30958, [BZ#1920689](#), [BZ#1921179](#), [BZ#1921802](#), [BZ#1916412](#), [BZ#1921812](#), [BZ#1920604](#))

Rust Toolset rebased to version 1.47.0

Rust Toolset has been updated to version 1.47.0. Notable changes include:

- The compile-time evaluated functions **const fn** have been improved and can now use control flow features, for example **if**, **while**, and **match**.
- The new **#[track_caller]** annotation can now be put on functions. Panics from annotated functions report the caller as the source.
- The Rust Standard Library now generically implements traits for arrays of any length. Previously, many of the trait implementations for arrays were only filled for lengths between 0 and 32.

For detailed instructions regarding usage, see [Using Rust Toolset](#).

(BZ#1883839)

The Logging System Role now supports property-based filter on its outputs

With this update, property-based filters have been added to the files output, the forwards output, and the remote_files output of the Logging System Role. The feature is provided by underlying the **rsyslog** sub-role, and is configurable via the Logging RHEL System Role. As a result, users can benefit from the ability of filtering log messages by the properties, such as hostname, tag, and the message itself is useful to manage logs.

(BZ#1889492)

The Logging RHEL System Role now supports **rsyslog** behavior

With this enhancement, **rsyslog** receives the message from Red Hat Virtualization and forwards the message to the **elasticsearch**.

(BZ#1889893)

The **ubi8/pause** container image is now available

Podman now uses the **ubi8/pause** instead of the **k8s.gcr.io/pause** container image to hold the network namespace information of the pod.

(BZ#1690785)

Podman rebased to version 2.1

The Podman utility has been updated to version 2.1. Notable enhancements include:

- Changes:
 - Updated Podman to 2.2.1 (from 2.0.5), Buildah to 1.19 (from 1.15.1), Skopeo to 1.2.1 (from 1.1.1), Udicia to 0.2.3 (from 0.2.2), and CRIU to 3.15 (0.3.4)
 - Docker-compatible volume API endpoints (Create, Inspect, List, Remove, Prune) are now available
 - Added an API endpoint for generating systemd unit files for containers

- The **podman play kube** command now features support for setting CPU and Memory limits for containers
- The **podman play kube** command now supports persistent volumes claims using Podman named volumes
- The **podman play kube** command now supports Kubernetes configmaps via the **--configmap** option
- Experimental support for shortname aliasing has been added. This is not enabled by default, but can be turned on by setting the environment variable **CONTAINERS_SHORT_NAME_ALIASING** to on. For more information see [Container image short names in Podman](#).
- The new **podman image command** has been added. This allows for an image to be mounted, read-only, to inspect its contents without creating a container from it.
- The **podman save** and **podman load** commands can now create and load archives containing multiple images.
- Podman will now retry pulling an image at most 3 times if a pull fails due to network errors.
- Bug Fixes:
 - Fixed a bug where running systemd in a container on a cgroups v1 system would fail.

The Buildah tool has been updated to version 1.19. Notable enhancements include:

- Changes:
 - The **buildah inspect** command supports inspecting manifests
 - The **buildah push** command supports pushing manifests lists and digests
 - Added support for **--manifest** flags
 - The **--arch** and **--os** and **--variant** options has been added to select architecture and OS
 - Allow users to specify stdin into containers
 - Allow **FROM** to be overridden with **--from** option
 - Added **--ignorefile** flag to use alternate **.dockerignore** flags
 - short-names aliasing
 - Added **--policy** option to **buildah pull** command
 - Fix **buildah mount** command to display container names not IDs
 - Improved buildah completions
 - Use **--timestamp** rather than **--omit-timestamp** flag
 - Use pipes for copying
 - Added **--omit-timestamp** flag to **buildah bud** command

- Add VFS additional image store to container
- Allow "readonly" as alias to "ro" in mount options
- buildah, bud: support **--jobs=N** option for parallel execution

The Skopeo tool has been updated to version 1.2.1. Notable enhancements include:

- Changes:
 - Add multi-arch builds for upstream and stable skopeo image via Travis
 - Added support for digests in sync
 - Added **--all** sync flag to emulate copy **--all**
 - Added **--format** option to **skopeo inspect** command

The Udica tool has been updated to version 0.2.3. Notable enhancements include:

- Changes:
 - Enable container port, not the host port
 - Add **--version** option

The CRIU tool has been updated to version 3.15. Notable enhancements include:

- Changes:
 - Initial cgroup2 support
 - Legalized swrk API and add the ability for inheriting fds via it
 - External bind mounts and tasks-to-cgroups bindings
 - **ibcriu.so** (RPC wrapper) and plugins

(JIRA:RHELPLAN-55998)

CHAPTER 5. RHEL 8.3.0 RELEASE

5.1. NEW FEATURES

This part describes new features and major enhancements introduced in Red Hat Enterprise Linux 8.3.

5.1.1. Installer and image creation

Anaconda rebased to version 33.16

With this release, Anaconda has been rebased to version 33.16. This version provides the following notable enhancements over the previous version.

- The Installation Program now displays static IPv6 addresses on multiple lines and no longer resizes the windows.
- The Installation Program now displays supported NVDIMM device sector sizes.
- Host name is now configured correctly on an installed system having IPv6 static configuration.
- You can now use non-ASCII characters in disk encryption passphrase.
- The Installation Program displays a proper recommendation to create a new file system on /boot, /tmp, and all /var and /usr mount points except /usr/local and /var/www.
- The Installation Program now correctly checks the keyboard layout and does not change the status of the Keyboard Layout screen when the keyboard keys (ALT+SHIFT) are used to switch between different layouts and languages.
- Rescue mode no longer fails on systems with existing RAID1 partitions.
- Changing of the LUKS version of the container is now available in the **Manual Partitioning** screen.
- The Installation Program successfully finishes the installation without the **btrfs-progs** package.
- The Installation Program now uses the default LUKS2 version for an encrypted container.
- The Installation Program no longer crashes when a Kickstart file places physical volumes (PVs) of a Logical volume group (VG) on an **ignoredisk** list.
- Introduces a new mount path **/mnt/sysroot** for system root. This path is used to mount / of the target system. Usually, the physical root and the system root are the same, so **/mnt/sysroot** is attached to the same file system as **/mnt/sysimage**. The only exceptions are rpm-ostree systems, where the system root changes based on the deployment. Then, **/mnt/sysroot** is attached to a subdirectory of **/mnt/sysimage**. It is recommended to use **/mnt/sysroot** for chroot.

([BZ#1691319](#), [BZ#1679893](#), [BZ#1684045](#), [BZ#1688478](#), [BZ#1700450](#), [BZ#1720145](#), [BZ#1723888](#), [BZ#1754977](#), [BZ#1755996](#), [BZ#1784360](#), [BZ#1796310](#), [BZ#1871680](#))

GUI changes in RHEL Installation Program

The RHEL Installation Program now includes the following user settings on the Installation Summary window:

- Root password
- User creation

With this change, you can now configure a root password and create a user account before you begin the installation. Previously, you configured a root password and created a user account after you began the installation process.

A root password is used to log in to the administrator (also known as superuser or root) account which is used for system administration tasks. The user name is used to log in from a command line; if you install a graphical environment, then your graphical login manager uses the full name. For more details, see [Performing a standard RHEL installation](#) document.

(JIRA:RHELPLAN-40469)

Image Builder backend **osbuild-composer** replaces **lorax-composer**

The **osbuild-composer** backend replaces **lorax-composer**. The new service provides REST APIs for image building. As a result, users can benefit from a more reliable backend and more predictable output images.

(BZ#1836211)

Image Builder **osbuild-composer** supports a set of image types

With the **osbuild-composer** backend replacement, the following set of image types supported in **osbuild-composer** this time:

- TAR Archive (.tar)
- QEMU QCOW2 (.qcow2)
- VMware Virtual Machine Disk (.vmdk)
- Amazon Machine Image (.ami)
- Azure Disk Image (.vhd)
- OpenStack Image (.qcow2)

The following outputs are not supported this time:

- ext4-filesystem
- partitioned-disk
- Alibaba Cloud
- Google GCE

(JIRA:RHELPLAN-42617)

Image Builder now supports push to clouds through GUI

With this enhancement, when creating images, users can choose the option of pushing to **Azure** and **AWS** service clouds through GUI **Image Builder**. As a result, users can benefit from easier uploads and instantiation.

(JIRA:RHELPLAN-30878)

5.1.2. RHEL for Edge

Introducing RHEL for Edge images

With this release, you can now create customized RHEL images for Edge servers.

You can use Image Builder to create RHEL for Edge images, and then use RHEL installer to deploy them on AMD and Intel 64-bit systems. Image Builder generates a RHEL for Edge image as **rhel-edge-commit** in a **.tar** file.

A RHEL for Edge image is an **rpm-ostree** image that includes system packages for remotely installing RHEL on Edge servers.

The system packages include:

- Base OS package
- Podman as the container engine

You can customize the image to configure the OS content as per your requirements, and can deploy them on physical and virtual machines.

With a RHEL for Edge image, you can achieve the following:

- Atomic upgrades, where the state of each update is known and no changes are seen until you reboot the device.
- Custom health checks using Greenboot and intelligent rollbacks for resiliency in case of failed upgrades.
- Container-focused workflows, where you can separate core OS updates from the application updates, and test and deploy different versions of applications.
- Optimized OTA payloads for low-bandwidth environments.
- Custom health checks using Greenboot to ensure resiliency.

For more information about composing, installing, and managing RHEL for Edge images, see [Composing, Installing, and Managing RHEL for Edge images](#) .

(JIRA:RHELPLAN-56676)

5.1.3. Software management

The default value for the **best** dnf configuration option has been changed from **True** to **False**

With this update, the value for the **best** dnf configuration option has been set to **True** in the default configuration file to retain the original dnf behavior. As a result, for users that use the default configuration file the behavior remains unchanged.

If you provide your own configuration files, make sure that the **best=True** option is present to retain the original behavior.

(BZ#1832869)

New **--norepopath** option for the **dnf reposync** command is now available

Previously, the **reposync** command created a subdirectory under the **--download-path** directory for each downloaded repository by default. With this update, the **--norepopath** option has been introduced, and **reposync** does not create the subdirectory. As a result, the repository is downloaded directly into the directory specified by **--download-path**. This option is also present in the **YUM v3**.

([BZ#1842285](#))

Ability to enable and disable the **libdnf** plugins

Previously, subscription checking was hardcoded into the RHEL version of the **libdnf** plug-ins. With this update, the **microdnf** utility can enable and disable the **libdnf** plug-ins, and subscription checking can now be disabled the same way as in DNF. To disable subscription checking, use the **--disableplugin=subscription-manager** command. To disable all plug-ins, use the **--noplugins** command.

([BZ#1781126](#))

5.1.4. Shells and command-line tools

ReaR updates

RHEL 8.3 introduces a number of updates to the Relax-and-Recover (**ReaR**) utility. Notable changes include:

- Support for the third-party Rubrik Cloud Data Management (CDM) as external backup software has been added. To use it, set the **BACKUP** option in the configuration file to **CDM**.
- Creation of a rescue image with a file larger than 4 GB on the IBM POWER, little endian architecture has been enabled.
- Disk layout created by **ReaR** no longer includes entries for Rancher 2 Longhorn iSCSI devices and file systems.

([BZ#1743303](#))

smartmontools rebased to version 7.1

The **smartmontools** package has been upgraded to version 7.1, which provides multiple bug fixes and enhancements. Notable changes include:

- HDD, SSD and USB additions to the drive database.
- New options **-j** and **--json** to enable JSON output mode.
- Workaround for the incomplete **Log** subpages response from some SAS SSDs.
- Improved handling of **READ CAPACITY** command.
- Various improvements for the decoding of the log pages.

([BZ#1671154](#))

opencryptoki rebased to version 3.14.0

The **opencryptoki** packages have been upgraded to version 3.14.0, which provides multiple bug fixes and enhancements. Notable changes include:

- EP11 cryptographic service enhancements:

- Dilithium support
- Edwards-curve digital signature algorithm (EdDSA) support
- Support of Rivest–Shamir–Adleman optimal asymmetric encryption padding (RSA-OAEP) with non-SHA1 hash and mask generation function (MGF)
- Enhanced process and thread locking
- Enhanced **btree** and object locking
- Support for new IBM Z hardware z15
- Support of multiple token instances for trusted platform module (TPM), IBM cryptographic architecture (ICA) and integrated cryptographic service facility (ICSF)
- Added a new tool **p11sak**, which lists the token keys in an **openCryptoki** token repository
- Added a utility to migrate a token repository to FIPS compliant encryption
- Fixed **pkcsep11_migrate** tool
- Minor fixes of the ICSF software

(BZ#1780293)

gpgme rebased to version 1.13.1.

The **gpgme** packages have been upgraded to upstream version 1.13.1. Notable changes include:

- New context flags **no-symkey-cache** (has an effect when used with GnuPG 2.2.7 or later), **request-origin** (has an effect when used with GnuPG 2.2.6 or later), **auto-key-locate**, and **trust-model** have been introduced.
- New tool **gpgme-json** as native messaging server for web browsers has been added. As of now, the public key encryption and decryption is supported.
- New encryption API to support direct key specification including hidden recipients option and taking keys from a file has been introduced. This also allows the use of a subkey.

(BZ#1829822)

5.1.5. Infrastructure services

powertop rebased to version 2.12

The **powertop** packages have been upgraded to version 2.12. Notable changes over the previously available version 2.11 include:

- Use of Device Interface Power Management (DIPM) for SATA link PM.
- Support for Intel Comet Lake mobile and desktop systems, the Skylake server, and the Atom-based Tremont architecture (Jasper Lake).

(BZ#1783110)

tuned rebased to version 2.14.0

The **tuned** packages have been upgraded to upstream version 2.14.0. Notable enhancements include:

- The **optimize-serial-console** profile has been introduced.
- Support for a post loaded profile has been added.
- The **irqbalance** plugin for handling **irqbalance** settings has been added.
- Architecture specific tuning for Marvell ThunderX and AMD based platforms has been added.
- Scheduler plugin has been extended to support **cgroups-v1** for CPU affinity setting.

([BZ#1792264](#))

tcpdump rebased to version 4.9.3

The **tcpdump** utility has been updated to version 4.9.3 to fix Common Vulnerabilities and Exposures (CVE).

([BZ#1804063](#))

libpcap rebased to version 1.9.1

The **libpcap** packages have been updated to version 1.9.1 to fix Common Vulnerabilities and Exposures (CVE).

([BZ#1806422](#))

iperf3 now supports sctp option on the client side

With this enhancement, the user can use Stream Control Transmission Protocol (SCTP) instead of Transmission Control Protocol (TCP) on the client side of testing network throughput.

The following options for **iperf3** are now available on the client side of testing:

- **--sctp**
- **--xbind**
- **--nstreams**

To obtain more information, see **Client Specific Options** in the **iperf3** man page.

([BZ#1665142](#))

iperf3 now supports SSL

With this enhancement, the user can use RSA authentication between the client and the server to restrict the connections to the server only to legitimate clients.

The following options for **iperf3** are now available on the server side:

- **--rsa-private-key-path**
- **--authorized-users-path**

The following options for **iperf3** are now available on the client side of communication:

- **--username**

- **--rsa-public-key-path**

([BZ#1700497](#))

bind rebased to 9.11.20

The **bind** package has been upgraded to version 9.11.20, which provides multiple bug fixes and enhancements. Notable changes include:

- Increased reliability on systems with many CPU cores by fixing several race conditions.
- Detailed error reporting: **dig** and other tools can now print the Extended DNS Error (EDE) option, if it is present.
- Message IDs in inbound DNS Zone Transfer Protocol (AXFR) transfers are checked and logged, when they are inconsistent.

([BZ#1818785](#))

A new optimize-serial-console TuneD profile to reduce I/O to serial consoles by lowering the printk value

With this update, a new **optimize-serial-console** TuneD profile is available. In some scenarios, kernel drivers can send large amounts of I/O operations to the serial console. Such behavior can cause temporary unresponsiveness while the I/O is written to the serial console. The **optimize-serial-console** profile reduces this I/O by lowering the **printk** value from the default of **7 4 1 7** to **4 4 1 7**. Users with a serial console who wish to make this change on their system can instrument their system as follows:

```
# tuned-adm profile throughput-performance optimize-serial-console
```

As a result, users will have a lower **printk** value that persists across a reboot, which reduces the likelihood of system hangs.

This TuneD profile reduces the amount of I/O written to the serial console by removing debugging information. If you need to collect this debugging information, you should ensure this profile is not enabled and that your **printk** value is set to **7 4 1 7**. To check the value of **printk** run:

```
# cat /proc/sys/kernel/printk
```

([BZ#1840689](#))

New TuneD profiles added for the AMD-based platforms

In RHEL 8.3, the **throughput-performance** TuneD profile was updated to include tuning for the AMD-based platforms. There is no need to change any parameter manually and the tuning is automatically applied on the **AMD** system. The **AMD Epyc Naples** and **Rome** systems alters the following parameters in the default **throughput-performance** profile:

sched_migration_cost_ns=5000000 and **kernel.numa_balancing=0**

With this enhancement, the system performance is improved by ~5%.

([BZ#1746957](#))

memcached rebased to version 1.5.22

The **memcached** packages have been upgraded to version 1.5.22. Notable changes over the previous version include:

- TLS has been enabled.
- The **-o inline_ascii_response** option has been removed.
- The **-Y [authfile]** option has been added along with authentication mode for the ASCII protocol.
- **memcached** can now recover its cache between restarts.
- New experimental meta commands have been added.
- Various performance improvements.

([BZ#1809536](#))

5.1.6. Security

Cyrus SASL now supports channel bindings with the **SASL/GSSAPI** and **SASL/GSS-SPNEGO** plug-ins

This update adds support for channel bindings with the **SASL/GSSAPI** and **SASL/GSS-SPNEGO** plug-ins. As a result, when used in the **openldap** libraries, this feature enables **Cyrus SASL** to maintain compatibility with and access to Microsoft Active Directory and Microsoft Windows systems which are introducing mandatory channel binding for LDAP connections.

([BZ#1817054](#))

Libreswan rebased to 3.32

With this update, Libreswan has been rebased to upstream version 3.32, which includes several new features and bug fixes. Notable features include:

- Libreswan no longer requires separate FIPS 140-2 certification.
- Libreswan now implements the cryptographic recommendations of RFC 8247, and changes the preference from SHA-1 and RSA-PKCS v1.5 to SHA-2 and RSA-PSS.
- Libreswan supports XFRMi virtual ipsecXX interfaces that simplify writing firewall rules.
- Recovery of crashed and rebooted nodes in a full-mesh encryption network is improved.

([BZ#1820206](#))

The **libssh** library has been rebased to version 0.9.4

The **libssh** library, which implements the SSH protocol, has been upgraded to version 0.9.4.

This update includes bug fixes and enhancements, including:

- Added support for **Ed25519** keys in PEM files.
- Added support for **diffie-hellman-group14-sha256** key exchange algorithm.
- Added support for **localuser** in **Match** keyword in the **libssh** client configuration file.

- **Match** criteria keyword arguments are now case-sensitive (note that keywords are case-insensitive, but keyword arguments are case-sensitive)
- Fixed CVE-2019-14889 and CVE-2020-1730.
- Added support for recursively creating missing directories found in the path string provided for the known hosts file.
- Added support for **OpenSSH** keys in PEM files with comments and leading white spaces.
- Removed the **OpenSSH** server configuration inclusion from the **libssh** server configuration.

([BZ#1804797](#))

gnutls rebased to 3.6.14

The **gnutls** packages have been rebased to upstream version 3.6.14. This version provides many bug fixes and enhancements, most notably:

- **gnutls** now rejects certificates with **Time** fields that contain invalid characters or formatting.
- **gnutls** now checks trusted CA certificates for minimum key sizes.
- When displaying an encrypted private key, the **certtool** utility no longer includes its plain text description.
- Servers using **gnutls** now advertise OCSP-stapling support.
- Clients using **gnutls** now send OCSP staples only on request.

([BZ#1789392](#))

gnutls FIPS DH checks now conform with NIST SP 800-56A rev. 3

This update of the **gnutls** packages provides checks required by NIST Special Publication 800-56A Revision 3, sections 5.7.1.1 and 5.7.1.2, step 2. The change is necessary for future FIPS 140-2 certifications. As a result, **gnutls** now accept only 2048-bit or larger parameters from RFC 7919 and RFC 3526 during the Diffie-Hellman key exchange when operating in FIPS mode.

([BZ#1849079](#))

gnutls now performs validations according to NIST SP 800-56A rev 3

This update of the **gnutls** packages adds checks required by NIST Special Publication 800-56A Revision 3, sections 5.6.2.2.2 and 5.6.2.1.3, step 2. The addition prepares **gnutls** for future FIPS 140-2 certifications. As a result, **gnutls** perform additional validation steps for generated and received public keys during the Diffie-Hellman key exchange when operating in FIPS mode.

([BZ#1855803](#))

update-crypto-policies and fips-mode-setup moved into crypto-policies-scripts

The **update-crypto-policies** and **fips-mode-setup** scripts, which were previously included in the **crypto-policies** package, are now moved into a separate RPM subpackage **crypto-policies-scripts**. The package is automatically installed through the Recommends dependency on regular installations. This enables the **ubi8/ubi-minimal** image to avoid the inclusion of the Python language interpreter and thus reduces the image size.

([BZ#1832743](#))

OpenSC rebased to version 0.20.0

The **opensc** package has been rebased to version 0.20.0 which addresses multiple bugs and security issues. Notable changes include:

- With this update, **CVE-2019-6502**, **CVE-2019-15946**, **CVE-2019-15945**, **CVE-2019-19480**, **CVE-2019-19481** and **CVE-2019-19479** security issues are fixed.
- The OpenSC module now supports the **C_WrapKey** and **C_UnwrapKey** functions.
- You can now use the facility to detect insertion and removal of card readers as expected.
- The **pkcs11-tool** utility now supports the **CKA_ALLOWED_MECHANISMS** attribute.
- This update allows default detection of the **OsEID** cards.
- The OpenPGP Card v3 now supports **Elliptic Curve Cryptography** (ECC).
- The PKCS#11 URI now truncates the reader name with ellipsis.

([BZ#1810660](#))

stunnel rebased to version 5.56

With this update, the **stunnel** encryption wrapper has been rebased to upstream version 5.56, which includes several new features and bug fixes. Notable features include:

- New **ticketKeySecret** and **ticketMacSecret** options that control confidentiality and integrity protection of the issued session tickets. These options enable you to resume sessions on other nodes in a cluster.
- New **curves** option to control the list of elliptic curves in OpenSSL 1.1.0 and later.
- New **ciphersuites** option to control the list of permitted TLS 1.3 ciphersuites.
- Added **sslVersion**, **sslVersionMin** and **sslVersionMax** for OpenSSL 1.1.0 and later.

([BZ#1808365](#))

libkcapi rebased to version 1.2.0

The **libkcapi** package has been rebased to upstream version 1.2.0, which includes minor changes.

([BZ#1683123](#))

setools rebased to 4.3.0

The **setools** package, which is a collection of tools designed to facilitate SELinux policy analysis, has been upgraded to version 4.3.0.

This update includes bug fixes and enhancements, including:

- Revised **sediff** method for Type Enforcement (TE) rules, which significantly reduces memory and runtime issues.
- Added **infiniband** context support to **seinfo**, **sediff**, and **apol**.
- Added **apol** configuration for the location of the Qt assistant tool used to display online documentation.

- Fixed **sediff** issues with:
 - Properties header displaying when not requested.
 - Name comparison of **type_transition** files.
- Fixed permission of map socket **sendto** information flow direction.
- Added methods to the **TypeAttribute** class to make it a complete Python collection.
- **Genfscon** now looks up classes, rather than using fixed values which were dropped from **libsepol**.

The **setools** package requires the following packages:

- **setools-console**
- **setools-console-analyses**
- **setools-gui**

([BZ#1820079](#))

Individual CephFS files and directories can now have SELinux labels

The Ceph File System (CephFS) has recently enabled storing SELinux labels in the extended attributes of files. Previously, all files in a CephFS volume were labeled with a single common label **system_u:object_r:cephfs_t:s0**. With this enhancement, you can change the labels for individual files, and SELinux defines the labels of newly created files based on transition rules. Note that previously unlabeled files still have the **system_u:object_r:cephfs_t:s0** label until explicitly changed.

([BZ#1823764](#))

OpenSCAP rebased to version 1.3.3

The **openscap** packages have been upgraded to upstream version 1.3.3, which provides many bug fixes and enhancements over the previous version, most notably:

- Added the **autotailor** script that enables you to generate tailoring files using a command-line interface (CLI).
- Added the timezone part to the Extensible Configuration Checklist Description Format (XCCDF) TestResult start and end time stamps
- Added the **yamlfilecontent** independent probe as a draft implementation.
- Introduced the **urn:xccdf:fix:script:kubernetes** fix type in XCCDF.
- Added ability to generate the **machineconfig** fix.
- The **oscap-podman** tool can now detect ambiguous scan targets.
- The **rpmverifyfile** probe can now verify files from the **/bin** directory.
- Fixed crashes when complicated regexes are executed in the **textfilecontent58** probe.
- Evaluation characteristics of the XCCDF report are now consistent with OVAL entities from the **system_info** probe.

- Fixed file-path pattern matching in offline mode in the **textfilecontent58** probe.
- Fixed infinite recursion in the **systemdunitdependency** probe.

([BZ#1829761](#))

SCAP Security Guide now provides a profile aligned with the CIS RHEL 8 Benchmark v1.0.0

With this update, the **scap-security-guide** packages provide a profile aligned with the CIS Red Hat Enterprise Linux 8 Benchmark v1.0.0. The profile enables you to harden the configuration of the system using the guidelines by the Center for Internet Security (CIS). As a result, you can configure and automate compliance of your RHEL 8 systems with CIS by using the CIS Ansible Playbook and the CIS SCAP profile.

Note that the **rpm_verify_permissions** rule in the CIS profile does not work correctly.

([BZ#1760734](#))

scap-security-guide now provides a profile that implements HIPAA

This update of the **scap-security-guide** packages adds the Health Insurance Portability and Accountability Act (HIPAA) profile to the RHEL 8 security compliance content. This profile implements recommendations outlined on the [The HIPAA Privacy Rule](#) website.

The HIPAA Security Rule establishes U.S. national standards to protect individuals' electronic personal health information that is created, received, used, or maintained by a covered entity. The Security Rule requires appropriate administrative, physical, and technical safeguards to ensure the confidentiality, integrity, and security of electronically protected health information.

([BZ#1832760](#))

scap-security-guide rebased to 0.1.50

The **scap-security-guide** packages, which contain the latest set of security policies for Linux systems, have been upgraded to version 0.1.50.

This update includes bug fixes and enhancements, most notably:

- Ansible content has been improved: numerous rules contain Ansible remediations for the first time and other rules have been updated to address bug fixes.
- Fixes and improvements to the **scap-security-guide** content for scanning RHEL7 systems, including:
 - The **scap-security-guide** packages now provide a profile aligned with the CIS RHEL 7 Benchmark v2.2.0. Note that the **rpm_verify_permissions** rule in the CIS profile does not work correctly; see the [rpm_verify_permissions fails in the CIS profile](#) known issue.
 - The SCAP Security Guide profiles now correctly disable and mask services that should not be started.
 - The **audit_rules_privileged_commands** rule in the **scap-security-guide** packages now works correctly for privileged commands.
 - Remediation of the **dconf_gnome_login_banner_text** rule in the **scap-security-guide** packages no longer incorrectly fails.

([BZ#1815007](#))

SCAP Workbench can now generate results-based remediations from tailored profiles

With this update, you can now generate result-based remediation roles from tailored profiles using the **SCAP Workbench** tool.

(BZ#1640715)

New Ansible role provides automated deployments of Clevis clients

This update of the **rhel-system-roles** package introduces the **nbde_client** RHEL system role. This Ansible role enables you to deploy multiple Clevis clients in an automated way.

(BZ#1716040)

New Ansible role can now set up a Tang server

With this enhancement, you can deploy and manage a Tang server as part of an automated disk encryption solution with the new **nbde_server** system role. The **nbde_server** Ansible role, which is included in the **rhel-system-roles** package, supports the following features:

- Rotating Tang keys
- Deploying and backing up Tang keys

For more information, see [Rotating Tang server keys](#).

(BZ#1716039)

clevis rebased to version 13

The **clevis** packages have been rebased to version 13, which provides multiple bug fixes and enhancements. Notable changes include:

- **clevis luks unlock** can be used in the device with a key file in the non-interactive mode.
- **clevis encrypt tpm2** parses the **pcr_ids** field if the input is given as a JSON array.
- The **clevis-luks-unbind(1)** man page no longer refers only to LUKS v1.
- **clevis luks bind** does not write to an inactive slot anymore, if the password given is incorrect.
- **clevis luks bind** now works while the system uses the non-English locale.
- Added support for **tpm2-tools** 4.x.

(BZ#1818780)

clevis luks edit enables you to edit a specific pin configuration

This update of the **clevis** packages introduces the new **clevis luks edit** subcommand that enables you to edit a specific pin configuration. For example, you can now change the URL address of a Tang server and the **pcr_ids** parameter in a TPM2 configuration. You can also add and remove new **sss** pins and change the threshold of an **sss** pin.

(BZ#1436735)

clevis luks bind -y now allows automated binding

With this enhancement, Clevis supports automated binding with the **-y** parameter. You can now use the **-y** option with the **clevis luks bind** command, which automatically answers subsequent prompts with yes. For example, when using a Tang pin, you are no longer required to manually trust Tang keys.

(BZ#1819767)

fapolicyd rebased to version 1.0

The **fapolicyd** packages have been rebased to version 1.0, which provides multiple bug fixes and enhancements. Notable changes include:

- The multiple thread synchronization problem has been resolved.
- Enhanced performance with reduced database size and loading time.
- A new trust option for the **fapolicyd** package in the **fapolicyd.conf** file has been added to customize trust back end. You can add all trusted files, binaries, and scripts to the new **/etc/fapolicyd/fapolicyd.trust** file.
- You can manage the **fapolicyd.trust** file using the CLI.
- You can clean or dump the database using the CLI.
- The **fapolicyd** package overrides the magic database for better decoding of scripts. The CLI prints MIME type of the file similar to the file command according to the override.
- The **/etc/fapolicyd/fapolicyd.rules** file supports a group of values as attribute values.
- The **fapolicyd** daemon has a **syslog_format** option for setting the format of the **audit/syslog** events.

(BZ#1817413)

fapolicyd now provides its own SELinux policy in fapolicyd-selinux

With this enhancement, the **fapolicyd** framework now provides its own SELinux security policy. The daemon is confined under the **fapolicyd_t** domain and the policy is installed through the **fapolicyd-selinux** subpackage.

(BZ#1714529)

USBGuard rebased to version 0.7.8

The **usbguard** packages have been rebased to version 0.7.8 which provides multiple bug fixes and enhancements. Notable changes include:

- The **HidePII=true|false** parameter in the **/etc/usbguard/usbguard-daemon.conf** file can now hide personally identifiable information from audit entries.
- The **AuthorizedDefault=keep|none|all|internal** parameter in the **/etc/usbguard/usbguard-daemon.conf** file can predefine authorization state of controller devices.
- With the new **with-connect-type** rule attribute, users can now distinguish the connection type of the device.
- Users can now append temporary rules with the **-t** option. Temporary rules remain in memory only until the daemon restarts.
- **usbguard list-rules** can now filter rules according to certain properties.

- **usbguard generate-policy** can now generate a policy for specific devices.
- The **usbguard allow|block|reject** command can now handle rule strings, and a target is applied on each device that matches the specified rule string.
- New subpackages **usbguard-notifier** and **usbguard-selinux** are included.

([BZ#1738590](#))

USBGuard provides many improvements for corporate desktop users

This addition to the USBGuard project contains enhancements and bug fixes to improve the usability for corporate desktop users. Important changes include:

- For keeping the **/etc/usbguard/rules.conf** rule file clean, users can define multiple configuration files inside the **RuleFolder=/etc/usbguard/rules.d/** directory. By default, the *RuleFolder* is specified in the **/etc/usbguard-daemon.conf** file.
- The **usbguard-notifier** tool now provides GUI notifications. The tool notifies the user whenever a device is plugged in or plugged out and whether the device is allowed, blocked, or rejected by any user.
- You can now include comments in the configuration files, because the **usbguard-daemon** no longer parses lines starting with **#**.

([BZ#1667395](#))

USBGuard now provides its own SELinux policy in **usbguard-selinux**

With this enhancement, the **USBGuard** framework now provides its own SELinux security policy. The daemon is confined under the **usbguard_t** domain and the policy is installed through the **usbguard-selinux** subpackage.

([BZ#1683567](#))

libcap now supports ambient capabilities

With this update, users are able to grant ambient capabilities at login and prevent the need to have root access for the appropriately configured processes.

([BZ#1487388](#))

The **libseccomp** library has been rebased to version 2.4.3

The **libseccomp** library, which provides an interface to the **seccomp** system call filtering mechanism, has been upgraded to version 2.4.3.

This update provides numerous bug fixes and enhancements. Notable changes include:

- Updated the **syscall** table for Linux v5.4-rc4.
- No longer defining **__NR_x** values for system calls that do not exist.
- **__SNR_x** is now used internally.
- Added **define** for **__SNR_ppoll**.
- Fixed a multiplexing issue with s390/s390x shm* system calls.

- Removed the **static** flag from the **libseccomp** tools compilation.
- Added support for **io-uring** related system calls.
- Fixed the Python module naming issue introduced in the v2.4.0 release; the module is named **seccomp** as it was previously.
- Fixed a potential memory leak identified by **clang** in the **scmp_bpf_sim** tool.

([BZ#1770693](#))

omamqp1 module is now supported

With this update, the **AMQP 1.0** protocol supports sending messages to a destination on the bus. Previously, Openstack used the **AMQP1** protocol as a communication standard, and this protocol can now log messages in AMQP messages. This update introduces the **rsyslog-omamqp1** sub-package to deliver the **omamqp1** output mode, which logs messages and sends them to the destination on the bus.

([BZ#1713427](#))

OpenSCAP compresses remote content

With this update, OpenSCAP uses **gzip** compression for transferring remote content. The most common type of remote content is text-based CVE feeds, which increase in size over time and typically have to be downloaded for every scan. The **gzip** compression reduces the bandwidth to 10% of bandwidth needed for uncompressed content. As a result, this reduces bandwidth requirements across the entire chain between the scanned system and the server that hosts the remote content.

([BZ#1855708](#))

SCAP Security Guide now provides a profile aligned with NIST-800-171

With this update, the **scap-security-guide** packages provide a profile aligned with the NIST-800-171 standard. The profile enables you to harden the system configuration in accordance with security requirements for protection of Controlled Unclassified Information (CUI) in non-federal information systems. As a result, you can more easily configure systems to be aligned with the NIST-800-171 standard.

([BZ#1762962](#))

5.1.7. Networking

The IPv4 and IPv6 connection tracking modules have been merged into the `nf_conntrack` module

This enhancement merges the **`nf_conntrack_ipv4`** and **`nf_conntrack_ipv6`** Netfilter connection tracking modules into the **`nf_conntrack`** kernel module. Due to this change, blacklisting the address family-specific modules no longer work in RHEL 8.3, and you can blacklist only the **`nf_conntrack`** module to disable connection tracking support for both the IPv4 and IPv6 protocols.

([BZ#1822085](#))

firewalld rebased to version 0.8.2

The **firewalld** packages have been upgraded to upstream version 0.8.2, which provides a number of bug fixes over the previous version. For details, see the [firewalld 0.8.2 Release Notes](#).

([BZ#1809636](#))

NetworkManager rebased to version 1.26.0

The **NetworkManager** packages have been upgraded to upstream version 1.26.0, which provides a number of enhancements and bug fixes over the previous version:

- NetworkManager resets the auto-negotiation, speed, and duplex setting to their original value when deactivating a device.
- Wi-Fi profiles connect now automatically if all previous activation attempts failed. This means that an initial failure to auto-connect to the network no longer blocks the automatism. A side effect is that existing Wi-Fi profiles that were previously blocked now connect automatically.
- The **nm-settings-nmcli(5)** and **nm-settings-dbus(5)** man pages have been added.
- Support for a number of bridge parameters has been added.
- Support for virtual routing and forwarding (VRF) interfaces has been added. For further details, see [Permanently reusing the same IP address on different interfaces](#).
- Support for Opportunistic Wireless Encryption mode (OWE) for Wi-Fi networks has been added.
- NetworkManager now supports 31-bit prefixes on IPv4 point-to-point links according to [RFC 3021](#).
- The **nmcli** utility now supports removing settings using the **nmcli connection modify <connection_name> remove <setting>** command.
- NetworkManager no longer creates and activates slave devices if a master device is missing.

For further information about notable changes, read the upstream release notes:

- [NetworkManager 1.26.0](#)
- [NetworkManager 1.24.0](#)

([BZ#1814746](#))

XDP is conditionally supported

Red Hat supports the eXpress Data Path (XDP) feature only if all of the following conditions apply:

- You load the XDP program on an AMD or Intel 64-bit architecture
- You use the **libxdp** library to load the program into the kernel
- The XDP program uses one of the following return codes: **XDP_ABORTED**, **XDP_DROP**, or **XDP_PASS**
- The XDP program does not use the XDP hardware offloading

For details about unsupported XDP features, see [Overview of XDP features that are available as Technology Preview](#)

([BZ#1889736](#))

xdp-tools is partially supported

The **xdp-tools** package, which contains user space support utilities for the kernel eXpress Data Path (XDP) feature, is now supported on the AMD and Intel 64-bit architectures. This includes the **libxdp**

library, the **xdp-loader** utility for loading XDP programs, and the **xdp-filter** example program for packet filtering. Note that the **xdpdump** utility for capturing packets from a network interface with XDP enabled is still a Technology Preview. (BZ#1820670)

The dracut utility by default now uses NetworkManager in initial RAM disk

Previously, the **dracut** utility was using a shell script to manage networking in the initial RAM disk, **initrd**. In certain cases, this could cause problems. For example, the NetworkManager sends another DHCP request, even if the script in the RAM disk has already requested an IP address, which could result in a timeout.

With this update, the **dracut** by default now uses the NetworkManager in the initial RAM disk and prevents the system from running into issues. In case you want to switch back to the previous implementation, and recreate the RAM disk images, use the following commands:

```
# echo 'add_dracutmodules+= " network-legacy "' > /etc/dracut.conf.d/enable-network-legacy.conf

# dracut -vf --regenerate-all
```

(BZ#1626348)

Network configuration in the kernel command line has been consolidated under the **ip** parameter

The **ipv6**, **netmask**, **gateway**, and **hostname** parameters to set the network configuration in the kernel command line have been consolidated under the **ip** parameter. The **ip** parameter accepts different formats, such as the following:

```
ip=__IP_address__:__peer__:__gateway_IP_address__:__net_mask__:__host_name__:__interface_name__:__configuration_method__
```

For further details about the individual fields and other formats this parameter accepts, see the description of the **ip** parameter in the **dracut.cmdline(7)** man page.

The **ipv6**, **netmask**, **gateway**, and **hostname** parameters are no longer available in RHEL 8.

(BZ#1905138)

5.1.8. Kernel

Kernel version in RHEL 8.3

Red Hat Enterprise Linux 8.3 is distributed with the kernel version 4.18.0-240.

(BZ#1839151)

Extended Berkeley Packet Filter for RHEL 8.3

The **Extended Berkeley Packet Filter (eBPF)** is an in-kernel virtual machine that allows code execution in the kernel space, in the restricted sandbox environment with access to a limited set of functions. The virtual machine executes a special assembly-like code.

The **eBPF** bytecode first loads to the kernel, followed by its verification, code translation to the native machine code with just-in-time compilation, and then the virtual machine executes the code.

Red Hat ships numerous components that utilize the **eBPF** virtual machine. Each component is in a different development phase, and thus not all components are currently fully supported. In RHEL 8.3, the following **eBPF** components are supported:

- The **BPF Compiler Collection (BCC)** tools package, which provides tools for I/O analysis, networking, and monitoring of Linux operating systems using **eBPF**
- The **BCC** library which allows the development of tools similar to those provided in the **BCC** tools package.
- The **eBPF for Traffic Control (tc)** feature, which enables programmable packet processing inside the kernel network data path.
- The **eXpress Data Path (XDP)** feature, which provides access to received packets before the kernel networking stack processes them, is supported under specific conditions. For more details, refer to the [Networking](#) section of Release Notes.
- The **libbpf** package, which is crucial for bpf related applications like **bpfftrace** and **bpf/xdp** development. For more details, refer to the dedicated release note [libbpf fully supported](#).
- The **xdp-tools** package, which contains userspace support utilities for the **XDP** feature, is now supported on the AMD and Intel 64-bit architectures. This includes the **libxdp** library, the **xdp-loader** utility for loading XDP programs, and the **xdp-filter** example program for packet filtering. Note that the **xdpdump** utility for capturing packets from a network interface with XDP enabled is still an unsupported Technology Preview. For more details, refer to the [Networking](#) section of Release Notes.

Note that all other **eBPF** components are available as Technology Preview, unless a specific component is indicated as supported.

The following notable **eBPF** components are currently available as Technology Preview:

- The **bpfftrace** tracing language
- The **AF_XDP** socket for connecting the **eXpress Data Path (XDP)** path to user space

For more information regarding the Technology Preview components, see [Technology Previews](#).

([BZ#1780124](#))

Cornelis Networks Omni-Path Architecture (OPA) Host Software

Omni-Path Architecture (OPA) host software is fully supported in Red Hat Enterprise Linux 8.3. OPA provides Host Fabric Interface (HFI) hardware with initialization and setup for high performance data transfers (high bandwidth, high message rate, low latency) between compute and I/O nodes in a clustered environment.

([BZ#1893174](#))

TSX is now disabled by default

Starting with RHEL 8.3, the kernel now has the **Intel® Transactional Synchronization Extensions (TSX)** technology disabled by default to improve the OS security. The change applies to those CPUs that support disabling **TSX**, including the 2nd Generation Intel® Xeon® Scalable Processors (formerly known as Cascade Lake with Intel® C620 Series Chipsets).

For users whose applications do not use **TSX**, the change removes the default performance penalty of the **TSX Asynchronous Abort (TAA)** mitigations on the 2nd Generation Intel® Xeon® Scalable Processors.

The change also aligns the RHEL kernel behavior with upstream, where **TSX** has been disabled by default since Linux 5.4.

To enable **TSX**, add the **tsx=on** parameter to the kernel command line.

(BZ#1828642)

RHEL 8.3 now supports the page owner tracking feature

With this update, you can use the page owner tracking feature to observe the kernel memory utilization at the page allocation level.

To enable the page tracker, execute the following steps :

```
# grubby --args="page_owner=on" --update-kernel=0
# reboot
```

As a result, the page owner tracker will track the kernel memory consumption, which helps to debug kernel memory leaks and detect the drivers that use a lot of memory.

(BZ#1825414)

EDAC for AMD EPYC™ 7003 Series Processors is now supported

This enhancement provides Error Detection And Correction (EDAC) device support for AMD EPYC™ 7003 Series Processors. Previously, corrected (CEs) and uncorrected (UEs) memory errors were not reported on systems based on AMD EPYC™ 7003 Series Processors. With this update, such errors will now be reported using EDAC.

(BZ#1735611)

Flamegraph is now supported with perf tool

With this update, the **perf** command line tool supports flamegraphs to create a graphical representation of the system's performance. The **perf** data is grouped together into samples with similar stack backtraces. As a result, this data is converted into a visual representation to allow easier identification of computationally intensive areas of code. To generate a flamegraph using the **perf** tool, execute the following commands:

```
$ perf script record flamegraph -F 99 -g -- stress --cpu 1 --vm-bytes 128M --timeout 10s
stress: info: [4461] dispatching hogs: 1 cpu, 0 io, 0 vm, 0 hdd
stress: info: [4461] successful run completed in 10s
[ perf record: Woken up 1 times to write data ]
[ perf record: Captured and wrote 0.060 MB perf.data (970 samples) ]
$ perf script report flamegraph
dumping data to flamegraph.html
```

Note : To generate flamegraphs, install the **js-d3-flame-graph** rpm.

(BZ#1281843)

/dev/random and **/dev/urandom** are now conditionally powered by the Kernel Crypto API
DRBG

In FIPS mode, the **/dev/random** and **/dev/urandom** pseudorandom number generators are powered by the Kernel Crypto API Deterministic Random Bit Generator (DRBG). Applications in FIPS mode use the mentioned devices as a FIPS-compliant noise source, therefore the devices have to employ FIPS-approved algorithms. To achieve this goal, necessary hooks have been added to the **/dev/random** driver. As a result, the hooks are enabled in the FIPS mode and cause **/dev/random** and **/dev/urandom** to connect to the Kernel Crypto API DRBG.

(BZ#1785660)

libbpf fully supported

The **libbpf** package, crucial for bpf related applications like **bpftrace** and **bpf/xdp** development, is now fully supported.

It is a mirror of bpf-next linux tree **bpf-next/tools/lib/bpf** directory plus its supporting header files. The version of the package reflects the version of the Application Binary Interface (ABI).

(BZ#1759154)

lshw utility now provides additional CPU information

With this enhancement, the List Hardware utility (*lshw*) displays more CPU information. The CPU **version** field now provides the family, model and stepping details of the system processors in numeric format as **version: <family>.<model>.<stepping>**.

(BZ#1794049)

kernel-rt source tree has been updated to the RHEL 8.3 tree

The **kernel-rt** sources have been updated to use the latest Red Hat Enterprise Linux kernel source tree. The real-time patch set has also been updated to the latest upstream version, v5.6.14-rt7. Both of these updates provide a number of bug fixes and enhancements.

(BZ#1818138, BZ#1818142)

tpm2-tools rebased to version 4.1.1

The **tpm2-tools** package has been upgraded to version 4.1.1, which provides a number of command additions, updates, and removals. For more details, see the [Updates to tpm2-tools package in RHEL8.3](#) solution.

(BZ#1789682)

The Mellanox ConnectX-6 Dx network adapter is now fully supported

This enhancement adds the PCI IDs of the Mellanox ConnectX-6 Dx network adapter to the **mlx5_core** driver. On hosts that use this adapter, RHEL loads the **mlx5_core** driver automatically. This feature, previously available as a technology preview, is now fully supported in RHEL 8.3.

(BZ#1782831)

mlxsw driver rebased to version 5.7

The **mlxsw driver** is upgraded to upstream version 5.7 and include following new features:

- The shared buffer occupancy feature, which provides buffer occupancy data.
- The packet drop feature, which enables monitoring the **layer 2**, **layer 3**, **tunnels** and **access control list** drops.

- Packet trap policers support.
- Default port priority configuration support using Link Layer Discovery Protocol (LLDP) agent.
- Enhanced Transmission Selection (ETS) and Token Bucket Filter (TBF) queuing discipline offloading support.
- RED queuing discipline **nodrop** mode is enabled to prevent early packet drops.
- Traffic class SKB editing action **skbedit** priority feature enables changing packets metadata and it complements with **pedit** Traffic Class Offloading (TOS).

(BZ#1821646)

The crash kernel now expands memory reserve for kdump

With this enhancement, the **crashkernel=auto** argument now reserves more memory on machines with 4GB to 64GB memory capacity. Previously, due to limited memory reserve, the crash kernel failed to capture the crash dump as the kernel space and user space memory expanded. As a consequence, the crash kernel experienced an out-of-memory (OOM) error. This update helps to reduce the OOM error occurrences in the described scenario and expands the memory capacity for **kdump** accordingly.

(BZ#1746644)

5.1.9. File systems and storage

LVM can now manage VDO volumes

LVM now supports the Virtual Data Optimizer (VDO) segment type. As a result, you can now use LVM utilities to create and manage VDO volumes as native LVM logical volumes.

VDO provides inline block-level deduplication, compression, and thin provisioning features.

For more information, see [Deduplicating and compressing logical volumes on RHEL](#).

(BZ#1598199)

The SCSI stack now works better with high-performance adapters

The performance of the SCSI stack has been improved. As a result, next-generation, high performance host bus adapters (HBAs) are now capable of higher IOPS (I/Os per second) on RHEL.

(BZ#1761928)

The megaraid_sas driver has been updated to the latest version

The **megaraid_sas** driver has been updated to version 07.713.01.00-rc1. This update provides several bug fixes and enhancements relating to improving performance, better stability of supported MegaRAID adapters, and a richer feature set.

(BZ#1791041)

Stratis now lists the pool name on error

When you attempt to create a Stratis pool on a block device that is already in use by an existing Stratis pool, the **stratis** utility now reports the name of the existing pool. Previously, the utility listed only the UUID label of the pool.

([BZ#1734496](#))

FPIN ELS frame notification support

The **lpfc** Fibre Channel (FC) driver now supports Fabric Performance Impact Notifications (FPINs) regarding link integrity, which help identify link level issues and allows the switch to choose a more reliable path.

([BZ#1796565](#))

New commands to debug LVM on-disk metadata

The **pvck** utility, which is available from the **lvm2** package, now provides low-level commands to debug or rescue LVM on-disk metadata on physical volumes:

- To extract metadata, use the **pvck --dump** command.
- To repair metadata, use the **pvck --repair** command.

For more information, see the **pvck(8)** man page.

([BZ#1541165](#))

LVM RAID supports DM integrity to prevent data loss due to corrupted data on a device

It is now possible to add Device Mapper (DM) integrity to an LVM RAID configuration to prevent data loss. The integrity layer detects data corruption on a device and alerts the RAID layer to fix the corrupted data across the LVM RAID.

While RAID prevents data loss due to device failure, adding integrity to an LVM RAID array prevents data loss due to corrupted data on a device. You can add the integrity layer when you create a new LVM RAID, or you can add it to an LVM RAID that already exists.

([JIRA:RHELPLAN-39320](#))

Resilient Storage (GFS2) supported on AWS, Azure, and Aliyun public clouds

Resilient Storage (GFS2) is now supported on three major public clouds, Amazon (AWS), Microsoft (Azure) and Alibaba (Aliyun) with the introduction of shared block device support on those platforms. As a result GFS2 is now a true hybrid cloud cluster filesystem with options to use both on premises and in the public cloud. For information on configuring shared block storage on Microsoft Azure and on AWS, see [Deploying Red Hat Enterprise Linux 8 on public cloud platforms](#) . For information on configuring shared block storage on Alibaba Cloud, see [Configuring Shared Block Storage for a Red Hat High Availability Cluster on Alibaba Cloud](#).

([BZ#1900019](#))

Userspace now supports the latest **nfsdclld** daemon

Userspace now supports the latest **nfsdclld** daemon, which is the only namespace-aware client tracking method. This enhancement ensures client open or lock recovery from the containerized **knfsd** daemon without any data corruption.

([BZ#1817756](#))

nconnect now supports multiple concurrent connections

With this enhancement, you can use the **nconnect** functionality to create multiple concurrent connections to an NFS server, allowing for a different load balancing ability. Enable the **nconnect**

functionality with the **nconnect=X** NFS mount option, where *X* is the number of concurrent connections to use. The current limit is 16.

(BZ#1683394, BZ#1761352)

nfsdclld daemon for client information tracking is now supported

With this enhancement, the **nfsdclld** daemon is now the default method in tracking per-client information on a stable storage. As a result, the NFS v4 running in containers allows the clients to reclaim the opens or locks after a server restart.

(BZ#1817752)

5.1.10. High availability and clusters

pacemaker rebased to version 2.0.4

The Pacemaker cluster resource manager has been upgraded to upstream version 2.0.4, which provides a number of bug fixes.

(BZ#1828488)

New priority-fencing-delay cluster property

Pacemaker now supports the new **priority-fencing-delay** cluster property, which allows you to configure a two-node cluster so that in a split-brain situation the node with the fewest resources running is the node that gets fenced.

The **priority-fencing-delay** property can be set to a time duration. The default value for this property is 0 (disabled). If this property is set to a non-zero value, and the **priority** meta-attribute is configured for at least one resource, then in a split-brain situation the node with the highest combined priority of all resources running on it will be more likely to survive.

For example, if you set **pcs resource defaults priority=1** and **pcs property set priority-fencing-delay=15s** and no other priorities are set, then the node running the most resources will be more likely to survive because the other node will wait 15 seconds before initiating fencing. If a particular resource is more important than the rest, you can give it a higher priority.

The node running the master role of a promotable clone will get an extra 1 point if a priority has been configured for that clone.

Any delay set with **priority-fencing-delay** will be added to any delay from the **pcmk_delay_base** and **pcmk_delay_max** fence device properties. This behavior allows some delay when both nodes have equal priority, or both nodes need to be fenced for some reason other than node loss (for example, **on-fail=fencing** is set for a resource monitor operation). If used in combination, it is recommended that you set the **priority-fencing-delay** property to a value that is significantly greater than the maximum delay from **pcmk_delay_base** and **pcmk_delay_max**, to be sure the prioritized node is preferred (twice the value would be completely safe).

(BZ#1784601)

New commands for managing multiple sets of resource and operation defaults

It is now possible to create, list, change and delete multiple sets of resource and operation defaults. When you create a set of default values, you can specify a rule that contains **resource** and **op** expressions. This allows you, for example, to configure a default resource value for all resources of a particular type. Commands that list existing default values now include multiple sets of defaults in their output.

- The **pcs resource [op] defaults set create** command creates a new set of default values. When specifying rules with this command, only **resource** and **op** expressions, including **and**, **or** and parentheses, are allowed.
- The **pcs resource [op] defaults set delete | remove** command removes sets of default values.
- The **pcs resource [op] defaults set update** command changes the default values in a set.

(BZ#1817547)

Support for tagging cluster resources

It is now possible to tag cluster resources in a Pacemaker cluster with the **pcs tag** command. This feature allows you to administer a specified set of resources with a single command. You can also use the **pcs tag** command to remove or modify a resource tag, and to display the tag configuration.

The **pcs resource enable**, **pcs resource disable**, **pcs resource manage**, and **pcs resource unmanage** commands accept tag IDs as arguments.

(BZ#1684676)

Pacemaker now supports recovery by demoting a promoted resource rather than fully stopping it

It is now possible to configure a promotable resource in a Pacemaker cluster so that when a promote or monitor action fails for that resource, or the partition in which the resource is running loses quorum, the resource will be demoted but will not be fully stopped.

This feature can be useful when you would prefer that the resource continue to be available in the unpromoted mode. For example, if a database master's partition loses quorum, you might prefer that the database resource lose the **Master** role, but stay alive in read-only mode so applications that only need to read can continue to work despite the lost quorum. This feature can also be useful when a successful demote is both sufficient for recovery and much faster than a full restart.

To support this feature:

- The **on-fail** operation meta-attribute now accepts a **demote** value when used with **promote** actions, as in the following example:

```
pcs resource op add my-rsc promote on-fail="demote"
```

- The **on-fail** operation meta-attribute now accepts a **demote** value when used with **monitor** actions with both **interval** set to a nonzero value and **role** set to **Master**, as in the following example:

```
pcs resource op add my-rsc monitor interval="10s" on-fail="demote" role="Master"
```

- The **no-quorum-policy** cluster property now accepts a **demote** value. When set, if a cluster partition loses quorum, any promoted resources will be demoted but left running and all other resources will be stopped.

Specifying a **demote** meta-attribute for an operation does not affect how promotion of a resource is determined. If the affected node still has the highest promotion score, it will be selected to be promoted again.

(BZ#1837747, BZ#1843079)

New **SBD_SYNC_RESOURCE_STARTUP** SBD configuration parameter to improve synchronization with Pacemaker

To better control synchronization between SBD and Pacemaker, the **/etc/sysconfig/sbd** file now supports the **SBD_SYNC_RESOURCE_STARTUP** parameter. When Pacemaker and SBD packages from RHEL 8.3 or later are installed and SBD is configured with **SBD_SYNC_RESOURCE_STARTUP=true**, SBD contacts the Pacemaker daemon for information about the daemon's state.

In this configuration, the Pacemaker daemon will wait until it has been contacted by SBD, both before starting its subdaemons and before final exit. As a result, Pacemaker will not run resources if SBD cannot actively communicate with it, and Pacemaker will not exit until it has reported a graceful shutdown to SBD. This prevents the unlikely situation that might occur during a graceful shutdown when SBD fails to detect the brief moment when no resources are running before Pacemaker finally disconnects, which would trigger an unneeded reboot. Detecting a graceful shutdown using a defined handshake works in maintenance mode as well. The previous method of detecting a graceful shutdown on the basis of no running resources left had to be disabled in maintenance mode since running resources would not be touched on shutdown.

In addition, enabling this feature avoids the risk of a split-brain situation in a cluster when SBD and Pacemaker both start successfully but SBD is unable to contact pacemaker. This could happen, for example, due to SELinux policies. In this situation, Pacemaker would assume that SBD is functioning when it is not. With this new feature enabled, Pacemaker will not complete startup until SBD has contacted it. Another advantage of this new feature is that when it is enabled SBD will contact Pacemaker repeatedly, using a heartbeat, and it is able to panic the node if Pacemaker stops responding at any time.



NOTE

If you have edited your **/etc/sysconfig/sbd** file or configured SBD through PCS, then an RPM upgrade will not pull in the new **SBD_SYNC_RESOURCE_STARTUP** parameter. In these cases, to implement this feature you must manually add it from the **/etc/sysconfig/sbd.rpmnew** file or follow the procedure described in the **Configuration via environment** section of the **sbd(8)** man page.

([BZ#1718324](#), [BZ#1743726](#))

5.1.11. Dynamic programming languages, web and database servers

A new module stream: **ruby:2.7**

RHEL 8.3 introduces Ruby 2.7.1 in a new **ruby:2.7** module stream. This version provides a number of performance improvements, bug and security fixes, and new features over Ruby 2.6 distributed with RHEL 8.1.

Notable enhancements include:

- A new Compaction Garbage Collector (GC) has been introduced. This GC can defragment a fragmented memory space.
- Ruby yet Another Compiler-Compiler (Racc) now provides a command-line interface for the one-token Look-Ahead Left-to-Right – LALR(1) – parser generator.
- Interactive Ruby Shell (**irb**), the bundled Read-Eval-Print Loop (REPL) environment, now supports multi-line editing.

- Pattern matching, frequently used in functional programming languages, has been introduced as an experimental feature.
- Numbered parameter as the default block parameter has been introduced as an experimental feature.

The following performance improvements have been implemented:

- Fiber cache strategy has been changed to accelerate fiber creation.
- Performance of the **CGI.escapeHTML** method has been improved.
- Performance of the **Monitor** class and **MonitorMixin** module has been improved.

In addition, automatic conversion of keyword arguments and positional arguments has been deprecated. In Ruby 3.0, positional arguments and keyword arguments will be separated. For more information, see the [upstream documentation](#).

To suppress warnings against experimental features, use the **-W:no-experimental** command-line option. To disable a deprecation warning, use the **-W:no-deprecated** command-line option or add **Warning[:deprecated] = false** to your code.

To install the **ruby:2.7** module stream, use:

```
# yum module install ruby:2.7
```

If you want to upgrade from the **ruby:2.6** stream, see [Switching to a later stream](#).

(BZ#1817135)

A new module stream: **nodejs:14**

A new module stream, **nodejs:14**, is now available. **Node.js 14**, included in RHEL 8.3, provides numerous new features and bug and security fixes over **Node.js 12** distributed in RHEL 8.1.

Notable changes include:

- The V8 engine has been upgraded to version 8.3.
- A new experimental WebAssembly System Interface (WASI) has been implemented.
- A new experimental Async Local Storage API has been introduced.
- The diagnostic report feature is now stable.
- The streams APIs have been hardened.
- Experimental modules warnings have been removed.

With the release of the [RHEA-2020:5101](#) advisory, RHEL 8 provides **Node.js 14.15.0**, which is the most recent Long Term Support (LTS) version with improved stability.

To install the **nodejs:14** module stream, use:

```
# yum module install nodejs:14
```

If you want to upgrade from the **nodejs:12** stream, see [Switching to a later stream](#).

([BZ#1815402](#), [BZ#1891809](#))

git rebased to version 2.27

The **git** packages have been upgraded to upstream version 2.27. Notable changes over the previously available version 2.18 include:

- The **git checkout** command has been split into two separate commands:
 - **git switch** for managing branches
 - **git restore** for managing changes within the directory tree
- The behavior of the **git rebase** command is now based on the **merge** workflow by default rather than the previous **patch+apply** workflow. To preserve the previous behavior, set the **rebase.backend** configuration variable to **apply**.
- The **git difftool** command can now be used also outside a repository.
- Four new configuration variables, **{author,committer}.{name,email}**, have been introduced to override **user.{name,email}** in more specific cases.
- Several new options have been added that enable users to configure SSL for communication with proxies.
- Handling of commits with log messages in non-UTF-8 character encoding has been improved in the **git fast-export** and **git fast-import** utilities.
- The **lfs** extension has been added as a new **git-lfs** package. Git Large File Storage (LFS) replaces large files with text pointers inside **Git** and stores the file contents on a remote server.

([BZ#1825114](#), [BZ#1783391](#))

Changes in Python

RHEL 8.3 introduces the following changes to the **python38:3.8** module stream:

- The **Python** interpreter has been updated to version 3.8.3, which provides several bug fixes.
- The **python38-pip** package has been updated to version 19.3.1, and **pip** now supports installing **manylinux2014** wheels.

Performance of the **Python 3.6** interpreter, provided by the **python3** packages, has been significantly improved.

The **ubi8/python-27**, **ubi8/python-36**, and **ubi8/python-38** container images now support installing the **pipenv** utility from a custom package index or a PyPI mirror if provided by the customer. Previously, **pipenv** could only be downloaded from the upstream PyPI repository, and if the upstream repository was unavailable, the installation failed.

([BZ#1847416](#), [BZ#1724996](#), [BZ#1827623](#), [BZ#1841001](#))

A new module stream: php:7.4

RHEL 8.3 introduces **PHP 7.4**, which provides a number of bug fixes and enhancements over version 7.3.

This release introduces a new experimental extension, Foreign Function Interface (FFI), which enables you to call native functions, access native variables, and create and access data structures defined in C libraries. The FFI extension is available in the **php-ffi** package.

The following extensions have been removed:

- The **wddx** extension, removed from **php-xml** package
- The **recode** extension, removed from the **php-recode** package.

To install the **php:7.4** module stream, use:

```
# yum module install php:7.4
```

If you want to upgrade from the **php:7.3** stream, see [Switching to a later stream](#).

For details regarding PHP usage on RHEL 8, see [Using the PHP scripting language](#).

([BZ#1797661](#))

A new module stream: **nginx:1.18**

The **nginx 1.18** web and proxy server, which provides a number of bug fixes, security fixes, new features and enhancements over version 1.16, is now available. Notable changes include:

- Enhancements to HTTP request rate and connection limiting have been implemented. For example, the **limit_rate** and **limit_rate_after** directives now support variables, including new **\$limit_req_status** and **\$limit_conn_status** variables. In addition, dry-run mode has been added for the **limit_conn_dry_run** and **limit_req_dry_run** directives.
- A new **auth_delay** directive has been added, which enables delayed processing of unauthorized requests.
- The following directives now support variables: **grpc_pass**, **proxy_upload_rate**, and **proxy_download_rate**.
- Additional PROXY protocol variables have been added, namely **\$proxy_protocol_server_addr** and **\$proxy_protocol_server_port**.

To install the **nginx:1.18** stream, use:

```
# yum module install nginx:1.18
```

If you want to upgrade from the **nginx:1.16** stream, see [Switching to a later stream](#).

([BZ#1826632](#))

A new module stream: **perl:5.30**

RHEL 8.3 introduces **Perl 5.30**, which provides a number of bug fixes and enhancements over the previously released **Perl 5.26**. The new version also deprecates or removes certain language features. Notable changes with significant impact include:

- The **Math::BigInt::CalcEmu**, **arybase**, and **B::Debug** modules have been removed
- File descriptors are now opened with a **close-on-exec** flag

- Opening the same symbol as a file and as a directory handle is no longer allowed
- Subroutine attributes now must precede subroutine signatures
- The **:locked** and **:uniq** attributes have been removed
- Comma-less variable lists in formats are no longer allowed
- A bare **<<** here-document operator is no longer allowed
- Certain formerly deprecated uses of an unescaped left brace (**{**) character in regular expression patterns are no longer permitted
- The **AUTOLOAD()** subroutine can no longer be inherited to non-method functions
- The **sort** pragma no longer allows specifying a **sort** algorithm
- The **B::OP::terse()** subroutine has been replaced by the **B::Concise::b_terse()** subroutine
- The **File::Glob::glob()** function has been replaced by the **File::Glob::bsd_glob()** function
- The **dump()** function now must be invoked fully qualified as **CORE::dump()**
- The yada-yada operator (**...**) is a statement now, it cannot be used as an expression
- Assigning a non-zero value to the **\$[** variable now returns a fatal error
- The **\$*** and **\$#** variables are no longer allowed
- Declaring variables using the **my()** function in a false condition branch is no longer allowed
- Using the **sysread()** and **syswrite()** functions on the **:utf8** handles now returns a fatal error
- The **pack()** function no longer returns malformed UTF-8 format
- Unicode code points with a value greater than **IV_MAX** are no longer allowed
- Unicode 12.1 is now supported

To upgrade from an earlier **perl** module stream, see [Switching to a later stream](#).

Perl 5.30 is also available as an s2i-enabled **ubi8/perl-530** container image.

([BZ#1713592](#), [BZ#1732828](#))

A new module stream: **perl-libwww-perl:6.34**

RHEL 8.3 introduces a new **perl-libwww-perl:6.34** module stream, which provides the **perl-libwww-perl** package for all versions of **Perl** available in RHEL 8. The non-modular **perl-libwww-perl** package, available since RHEL 8.0, which cannot be used with other **Perl** streams than 5.26, has been obsoleted by the new default **perl-libwww-perl:6.34** stream.

([BZ#1781177](#))

A new module stream: **perl-IO-Socket-SSL:2.066**

A new **perl-IO-Socket-SSL:2.066** module stream is now available. This module provides the **perl-IO-Socket-SSL** and **perl-Net-SSLeay** packages and it is compatible with all **Perl** streams available in RHEL 8.

([BZ#1824222](#))

The **squid:4** module stream rebased to version 4.11

The **Squid** proxy server, provided by the **squid:4** module stream, has been upgraded from version 4.4 to version 4.11. This release provides multiple bug and security fixes, and various enhancements, such as new configuration options.

([BZ#1829467](#))

Changes in the **httpd:2.4** module stream

RHEL 8.3 introduces the following notable changes to the Apache HTTP Server, available through the **httpd:2.4** module stream:

- The **mod_http2** module rebased to version 1.15.7
- Configuration changes in the **H2Upgrade** and **H2Push** directives
- A new **H2Padding** configuration directive to control padding of the HTTP/2 payload frames
- Numerous bug fixes.

([BZ#1814236](#))

Support for logging to **journald** from the **CustomLog** directive in **httpd**

It is now possible to output access (transfer) logs to **journald** from the Apache HTTP Server by using a new option for the **CustomLog** directive.

The supported syntax is as follows:

```
CustomLog journald:priority format|nickname
```

where *priority* is any priority string up to **debug** as used in the [LogLevel directive](#).

For example, to log to **journald** using the the **combined** log format, use:

```
CustomLog journald:info combined
```

Note that when using this option, the server performance might be lower than when logging directly to flat files.

([BZ#1209162](#))

5.1.12. Compilers and development tools

.NET 5 is now available on RHEL

.NET 5 is available on Red Hat Enterprise Linux 7, Red Hat Enterprise Linux 8, and OpenShift Container Platform. .NET 5 includes new language versions: C# 9 and F# 5.0. Significant performance improvements were made in the base libraries, GC and JIT. .NET 5 has single file applications, which

allows you to distribute .NET applications as a single executable, with all dependencies included. UBI8 images for .NET 5 are available from Red Hat container registry and can be used with OpenShift.

To use .NET 5, install the **dotnet-sdk-5.0** package:

```
$ sudo dnf install -y dotnet-sdk-5.0
```

For more information, see the [.NET 5 documentation](#).

([BZ#1944677](#))

New GCC Toolset 10

GCC Toolset 10 is a compiler toolset that provides recent versions of development tools. It is available as an Application Stream in the form of a Software Collection in the **AppStream** repository.

The GCC compiler has been updated to version 10.2.1, which provides many bug fixes and enhancements that are available in upstream GCC.

The following tools and versions are provided by GCC Toolset 10:

Tool	Version
GCC	10.2.1
GDB	9.2
Valgrind	3.16.0
SystemTap	4.3
Dyninst	10.1.0
binutils	2.35
elfutils	0.180
dwz	0.12
make	4.2.1
strace	5.7
ltrace	0.7.91
annobin	9.29

To install GCC Toolset 10, run the following command as root:

```
# yum install gcc-toolset-10
```

To run a tool from GCC Toolset 10:

```
$ scl enable gcc-toolset-10 tool
```

To run a shell session where tool versions from GCC Toolset 10 override system versions of these tools:

```
$ scl enable gcc-toolset-10 bash
```

For more information, see [Using GCC Toolset](#).

The GCC Toolset 10 components are available in the two container images:

- **rhel8/gcc-toolset-10-toolchain**, which includes the GCC compiler, the GDB debugger, and the **make** automation tool.
- **rhel8/gcc-toolset-10-perftools**, which includes the performance monitoring tools, such as SystemTap and Valgrind.

To pull a container image, run the following command as root:

```
# podman pull registry.redhat.io/<image_name>
```

Note that only the GCC Toolset 10 container images are now supported. Container images of earlier GCC Toolset versions are deprecated.

For details regarding the container images, see [Using the GCC Toolset container images](#).

(BZ#1842656)

Rust Toolset rebased to version 1.45.2

Rust Toolset has been updated to version 1.45.2. Notable changes include:

- The subcommand **cargo tree** for viewing dependencies is now included in **cargo**.
- Casting from floating point values to integers now produces a clamped cast. Previously, when a truncated floating point value was out of range for the target integer type the result was undefined behaviour of the compiler. Non-finite floating point values led to undefined behaviour as well. With this enhancement, finite values are clamped either to the minimum or the maximum range of the integer. Positive and negative infinity values are by default clamped to the maximum and minimum integer respectively, Not-a-Number(NaN) values to zero.
- Function-like procedural macros in expressions, patterns, and statements are now extended and stabilized.

For detailed instructions regarding usage, see [Using Rust Toolset](#).

(BZ#1820593)

LLVM Toolset rebased to version 10.0.1

LLVM Toolset has been upgraded to version 10.0.1. With this update, the **clang-libs** packages no longer include individual component libraries. As a result, it is no longer possible to link applications against them. To link applications against the **clang** libraries, use the **libclang-cpp.so** package.

For more information, see [Using LLVM Toolset](#).

(BZ#1820587)

Go Toolset rebased to version 1.14.7

Go Toolset has been upgraded to version 1.14.7. Notable changes include:

- The Go module system is now fully supported.
- SSL version 3.0 (SSLv3) is no longer supported. Notable Delve debugger enhancements include:
 - The new command **examinemem** (or **x**) for examining raw memory
 - The new command **display** for printing values of an expression during each stop of the program
 - The new **--tty** flag for supplying a Teletypewriter (TTY) for the debugged program
 - The new coredump support for Arm64
 - The new ability to print goroutine labels
 - The release of the Debug Adapter Protocol (DAP) server
 - The improved output from **dlv trace** and **trace** REPL (read-eval-print-loop) commands

For more information on Go Toolset, see [Using Go Toolset](#).

For more information on Delve, see the upstream [Delve documentation](#).

(BZ#1820596)

SystemTap rebased to version 4.3

The SystemTap instrumentation tool has been updated to version 4.3, which provides multiple bug fixes and enhancements. Notable changes include:

- Userspace probes can be targeted by hexadecimal **buildid** from **readelf -n**. This alternative to a path name enables matching binaries to be probed under any name, and thus allows a single script to target a range of different versions. This feature works well in conjunction with the elfutils **debuginfod** server.
- Script functions can use probe **\$context** variables to access variables in the probed location, which allows the SystemTap scripts to use common logic to work with a variety of probes.
- The **stapbpf** program improvements, including try-catch statements, and error probes, have been made to enable proper error tolerance in scripts running on the BPF backend.

For further information about notable changes, read the [upstream release notes](#) before updating.

(BZ#1804319)

Valgrind rebased to version 3.16.0

The Valgrind executable code analysis tool has been updated to version 3.16.0, which provides a number of bug fixes and enhancements over the previous version:

- It is now possible to dynamically change the value of many command-line options while your program is running under Valgrind: through **vgdb**, through a **gdb** connected to the Valgrind gdbserver, or through program client requests. To get a list of dynamically changeable options,

run the **valgrind --help-dyn-options** command.

- For the Cachegrind (**cg_annotate**) and Callgrind (**callgrind_annotate**) tools the **--auto** and **--show-percs** options now default to **yes**.
- The Memcheck tool produces fewer false positive errors on optimized code. In particular, Memcheck now better handles the case when the compiler transformed an **A && B** check into **B && A**, where **B** could be undefined and **A** was false. Memcheck also better handles integer equality checks and non-equality checks on partially defined values.
- The experimental Stack and Global Array Checking tool (**exp-sgcheck**) has been removed. An alternative for detecting stack and global array overruns is using the AddressSanitizer (ASAN) facility of GCC, which requires you to rebuild your code with the **-fsanitize=address** option.

([BZ#1804324](#))

elfutils rebased to version 0.180

The **elfutils** package has been updated to version 0.180, which provides multiple bug fixes and enhancements. Notable changes include:

- Better support for debug info for code built with GCC LTO (link time optimization). The **eu-readelf** and **libdw** utilities now can read and handle **.gnu.debuglto_** sections, and correctly resolve file names for functions that are defined across CUs (compile units).
- The **eu-nm** utility now explicitly identifies weak objects as **V** and common symbols as **C**.
- The **debuginfod** server can now index **.deb** archives and has a generic extension to add other package archive formats using the **-Z EXT[=CMD]** option. For example **-Z '.tar.zst=zstdcat'** indicates that archives ending with the **.tar.zst** extension should be unpacked using the **zstdcat** utility.
- The **debuginfo-client** tool has several new helper functions, such as **debuginfod_set_user_data**, **debuginfod_get_user_data**, **debuginfod_get_url** and **debuginfod_add_http_header**. It also supports **file://** URLs now.

([BZ#1804321](#))

GDB now supports process record and replay on IBM z15

With this enhancement, the GNU Debugger (GDB) now supports process record and replay with most of the new instructions of the IBM z15 processor (previously known as arch13). Note that the following instructions are currently not supported: SORTL (sort lists), DFLTCC (deflate conversion call), KDSA (compute digital signature authentication).

([BZ#1659535](#))

Marvell ThunderX2 performance monitoring events have been updated in papi

With this enhancement, a number of performance events specific to ThunderX2, including uncore events, have been updated. As a result, developers can better investigate system performance on Marvell ThunderX2 systems.

([BZ#1726070](#))

The glibc math library is now optimized for IBM Z

With this enhancement, the **libm** math functions were optimized to improve performance on IBM Z machines. Notable changes include:

- improved rounding mode handling to avoid superfluous floating point control register sets and extracts
- exploitation of conversion between z196 integer and float

(BZ#1780204)

An additional libffi-specific temporary directory is available now

Previously on hardened systems, the system-wide temporary directories may not have had permissions suitable for use with the **libffi** library.

With this enhancement, system administrators can now set the **LIBFFI_TMPDIR** environment variable to point to a libffi-specific temporary directory with both **write** and **exec** mount or selinux permissions.

(BZ#1723951)

Improved performance of **strstr()** and **strcasestr()**

With this update, the performance of the **strstr()** and **strcasestr()** functions has been improved across several supported architectures. As a result, users now benefit from significantly better performance of all applications using string and memory manipulation routines.

(BZ#1821531)

glibc now handles loading of a truncated locale archive correctly

If the archive of system locales has been previously truncated, either due to a power outage during upgrade or a disk failure, a process could terminate unexpectedly when loading the archive. This enhancement adds additional consistency checks to the loading of the locale archive. As a result, processes are now able to detect archive truncation and fall back to either non-archive installed locales or the default POSIX locale.

(BZ#1784525)

GDB now supports **debuginfod**

With this enhancement, the GNU Debugger (GDB) can now download debug information packages from centralized servers on demand using the elfutils **debuginfod** client library.

(BZ#1838777)

pcp rebased to version 5.1.1-3

The **pcp** package has been upgraded to version 5.1.1-3. Notable changes include:

- Updated service units and improved **systemd** integration and reliability for all the PCP services. Improved archive log rotation and more timely compression. Archived discovery bug fixes in the **pmproxy** protocol.
- Improved **pcp-atop**, **pcp-dstat**, **pmrep**, and related monitor tools along with metric labels reporting in the **pmrep** and export tools.
- Improved **bpfftrace**, **OpenMetrics**, MMV, the Linux kernel agent, and other collection agents. New metric collectors for the **Open vSwitch** and **RabbitMQ** servers.

- New host discovery **pmfind systemd** service, which replaces the standalone **pmmgr** daemon.

([BZ#1792971](#))

grafana rebased to version 6.7.3

The **grafana** package has been upgraded to version 6.7.3. Notable changes include:

- Generic **OAuth** role mapping support
- A new logs panel
- Multi-line text display in the table panel
- A new currency and energy units

([BZ#1807323](#))

grafana-pcp rebased to version 2.0.2

The **grafana-pcp** package has been upgraded to version 2.0.2. Notable changes include:

- Supports the multidimensional **eBPF** maps to be graphed in the flamegraph.
- Removes an auto-completion cache in the query editor, so that the PCP metrics can appear dynamically.

([BZ#1807099](#))

A new rhel8/pcp container image

The **rhel8/pcp** container image is now available in the Red Hat Container Registry. The image contains the Performance Co-Pilot (PCP) toolkit, which includes preinstalled **pcp-zeroconf** package and the **OpenMetrics** PMDA.

([BZ#1497296](#))

A new rhel8/grafana container image

The **rhel8/grafana** container image is now available in the Red Hat Container Registry. Grafana is an open source utility with metrics dashboard, and graph editor for the **Graphite**, **Elasticsearch**, **OpenTSDB**, **Prometheus**, **InfluxDB**, and **PCP** monitoring tool.

([BZ#1823834](#))

5.1.13. Identity Management

IdM backup utility now checks for required replica roles

The **ipa-backup** utility now checks if all of the services used in the IdM cluster, such as a Certificate Authority (CA), Domain Name System (DNS), and Key Recovery Agent (KRA) are installed on the replica where you are running the backup. If the replica does not have all these services installed, the **ipa-backup** utility exits with a warning, because backups taken on that host would not be sufficient for a full cluster restoration.

For example, if your IdM deployment uses an integrated Certificate Authority (CA), a backup run on a non-CA replica will not capture CA data. Red Hat recommends verifying that the replica where you perform an **ipa-backup** has all of the IdM services used in the cluster installed.

For more information, see [Preparing for data loss with IdM backups](#) .

([BZ#1810154](#))

New password expiration notification tool

Expiring Password Notification (EPN), provided by the **ipa-client-epn** package, is a standalone tool you can use to build a list of Identity Management (IdM) users whose passwords are expiring soon.

IdM administrators can use EPN to:

- Display a list of affected users in JSON format, which is calculated at runtime
- Calculate how many emails will be sent for a given day or date range
- Send password expiration email notifications to users

Red Hat recommends launching EPN once a day from an IdM client or replica with the included **ipa-epn.timer systemd** timer.

([BZ#913799](#))

JSS now provides a FIPS-compliant SSLContext

Previously, Tomcat used the SSLEngine directive from the Java Cryptography Architecture (JCA) SSLContext class. The default SunJSSE implementation is not compliant with the Federal Information Processing Standard (FIPS), therefore PKI now provides a FIPS-compliant implementation via JSS.

([BZ#1821851](#))

Checking the overall health of your public key infrastructure is now available

With this update, the public key infrastructure (PKI) Healthcheck tool reports the health of the PKI subsystem to the Identity Management (IdM) Healthcheck tool, which was introduced in RHEL 8.1. Executing the IdM Healthcheck invokes the PKI Healthcheck, which collects and returns the health report of the PKI subsystem.

The **pki-healthcheck** tool is available on any deployed RHEL IdM server or replica. All the checks provided by **pki-healthcheck** are also integrated into the **ipa-healthcheck** tool. **ipa-healthcheck** can be installed separately from the **idm:DL1** module stream.

Note that **pki-healthcheck** can also work in a standalone Red Hat Certificate System (RHCS) infrastructure.

([BZ#1770322](#))

Support for RSA PSS

With this enhancement, PKI now supports the RSA PSS (Probabilistic Signature Scheme) signing algorithm.

To enable this feature, set the following line in the **pkispawn** script file for a given subsystem:
pki_use_pss_rsa_signing_algorithm=True

As a result, all existing default signing algorithms for this subsystem (specified in its **CS.cfg** configuration file) will use the corresponding PSS version. For example, **SHA256withRSA** becomes **SHA256withRSA/PSS**

([BZ#1824948](#))

Directory Server exports the private key and certificate to a private name space when the service starts

Directory Server uses OpenLDAP libraries for outgoing connections, such as replication agreements. Because these libraries cannot access the network security services (NSS) database directly, Directory Server extracts the private key and certificates from the NSS database on instances with TLS encryption support to enable the OpenLDAP libraries to establish encrypted connections. Previously, Directory Server extracted the private key and certificates to the directory set in the **nsslapd-certdir** parameter in the **cn=config** entry (default: `/etc/dirsrv/slapped-<instance_name>/`). As a consequence, Directory Server stored the **Server-Cert-Key.pem** and **Server-Cert.pem** in this directory. With this enhancement, Directory Server extracts the private key and certificate to a private name space that **systemd** mounts to the `/tmp/` directory. As a result, the security has been increased.

(BZ#1638875)

Directory Server can now turn an instance to read-only mode if the disk monitoring threshold is reached

This update adds the **nsslapd-disk-monitoring-readonly-on-threshold** parameter to the **cn=config** entry. If you enable this setting, Directory Server switches all databases to read-only if disk monitoring is enabled and the free disk space is lower than the value you configured in **nsslapd-disk-monitoring-threshold**. With **nsslapd-disk-monitoring-readonly-on-threshold** set to **on**, the databases cannot be modified until Directory Server successfully shuts down the instance. This can prevent data corruption.

(BZ#1728943)

samba rebased to version 4.12.3

The *samba* packages have been upgraded to upstream version 4.12.3, which provides a number of bug fixes and enhancements over the previous version:

- Built-in cryptography functions have been replaced with GnuTLS functions. This improves the server message block version 3 (SMB3) performance and copy speed significantly.
- The minimum runtime support is now Python 3.5.
- The **write cache size** parameter has been removed because the previous write cache concept could reduce the performance on memory-constrained systems.
- Support for authenticating connections using Kerberos tickets with DES encryption types has been removed.
- The **vfs_netatalk** virtual file system (VFS) module has been removed.
- The **ldap ssl ads** parameter is marked as deprecated and will be removed in a future Samba version. For information about how to alternatively encrypt LDAP traffic and further details, see the [samba: removal of "ldap ssl ads" smb.conf option](#) solution.
- By default, Samba on RHEL 8.3 no longer supports the deprecated RC4 cipher suite. If you run Samba as a domain member in an AD that still requires RC4 for Kerberos authentication, use the **update-crypto-policies --set DEFAULT:AD-SUPPORT** command to enable support for the RC4 encryption type.

Samba automatically updates its **tdb** database files when the **smbd**, **nmbd**, or **winbind** service starts. Back up the database files before starting Samba. Note that Red Hat does not support downgrading **tdb** database files.

For further information about notable changes, read the [upstream release notes](#) before updating.

([BZ#1817557](#))

cockpit-session-recording rebased to version 4

The **cockpit-session-recording** module has been rebased to version 4. This version provides following notable changes over the previous version:

- Updated parent id in the **metainfo** file.
- Updated package manifest.
- Fixed **rpmmacro** to resolve correct path on CentOS7.
- Handled byte-array encoded journal data.
- Moved code out of deprecated React lifecycle functions.

([BZ#1826516](#))

krb5 rebased to version 1.18.2

The **krb5** packages have been upgraded to upstream version 1.18.2. Notable fixes and enhancements include:

- Single- and triple-DES encryption types have been removed.
- Draft 9 PKINIT has been removed as it is not needed for any of the supported versions of Active Directory.
- NegoEx mechanism plug-ins are now supported.
- Hostname canonicalization fallback is now supported (**dns_canonicalize_hostname = fallback**).

([BZ#1802334](#))

IdM now supports new Ansible management modules

This update introduces several **ansible-freeipa** modules for automating common Identity Management (IdM) tasks using Ansible playbooks:

- The **config** module allows setting global configuration parameters within IdM.
- The **dnsconfig** module allows modifying global DNS configuration.
- The **dnsforwardzone** module allows adding and removing DNS forwarders from IdM.
- The **dnsrecord** allows the management of DNS records. In contrast to the upstream **ipa_dnsrecord**, it allows multiple record management in one execution, and it supports more record types.
- The **dnszone** module allows configuring zones in the DNS server.
- The **service** module allows ensuring the presence and absence of services.
- The **vault** module allows ensuring the presence and absence of vaults and of the members of vaults.

Note that the **ipagroup** and **ipahostgroup** modules have been extended to include user and host group

membership managers, respectively. A group membership manager is a user or a group that can add members to a group or remove members from a group. For more information, see the **Variables** sections of the respective `/usr/share/doc/ansible-freeipa/README-*` files.

(JIRA:RHELPLAN-49954)

IdM now supports a new Ansible system role for certificate management

Identity Management (IdM) supports a new Ansible system role for automating certificate management tasks. The new role includes the following benefits:

- The role helps automate the issuance and renewal of certificates.
- The role can be configured to have the **ipa** certificate authority issue your certificates. In this way, you can use your existing IdM infrastructure to manage the certificate trust chain.
- The role allows you to specify the commands to be executed before and after a certificate is issued, for example the stopping and starting of services.

(JIRA:RHELPLAN-50002)

Identity Management now supports FIPS

With this enhancement, you can now use encryption types that are approved by the Federal Information Processing Standard (FIPS) with the authentication mechanisms in Identity Management (IdM). Note that a cross-forest trust between IdM and Active Directory is not FIPS compliant.

Customers who require FIPS but do not require an AD trust can now install IdM in FIPS mode.

(JIRA:RHELPLAN-43531)

OpenDNSSEC in **idm:DL1** rebased to version 2.1

The OpenDNSSEC component of the **idm:DL1** module stream has been upgraded to the 2.1 version series, which is the current long term upstream support version. OpenDNSSEC is an open source project driving the adoption of Domain Name System Security Extensions (DNSSEC) to further enhance Internet security. OpenDNSSEC 2.1 provides a number of bug fixes and enhancements over the previous version. For more information, read the upstream release notes:

<https://www.opendnssec.org/archive/releases/>

(JIRA:RHELPLAN-48838)

IdM now supports the deprecated RC4 cipher suite with a new system-wide cryptographic subpolicy

This update introduces the new **AD-SUPPORT** cryptographic subpolicy that enables the Rivest Cipher 4 (RC4) cipher suite in Identity Management (IdM).

As an administrator in the context of IdM-Active Directory (AD) cross-forest trusts, you can activate the new **AD-SUPPORT** subpolicy when AD is not configured to use Advanced Encryption Standard (AES). More specifically, Red Hat recommends enabling the new subpolicy if one of the following conditions applies:

- The user or service accounts in AD have RC4 encryption keys and lack AES encryption keys.
- The trust links between individual Active Directory domains have RC4 encryption keys and lack AES encryption keys.

To enable the **AD-SUPPORT** subpolicy in addition to the **DEFAULT** cryptographic policy, enter:

```
# update-crypto-policies --set DEFAULT:AD-SUPPORT
```

Alternatively, to upgrade trusts between AD domains in an AD forest so that they support strong AES encryption types, see the following Microsoft article: [AD DS: Security: Kerberos "Unsupported etype" error when accessing a resource in a trusted domain](#).

(BZ#1851139)

Adjusting to new Microsoft LDAP channel binding and LDAP signing requirements

With recent Microsoft updates, Active Directory (AD) flags the clients that do not use the default Windows settings for LDAP channel binding and LDAP signing. As a consequence, RHEL systems that use the System Security Services Daemon (SSSD) for direct or indirect integration with AD might trigger error Event IDs in AD upon successful Simple Authentication and Security Layer (SASL) operations that use the Generic Security Services Application Program Interface (GSSAPI).

To prevent these notifications, configure client applications to use the Simple and Protected GSSAPI Negotiation Mechanism (GSS-SPNEGO) SASL mechanism instead of GSSAPI. To configure SSSD, set the **ldap_sasl_mech** option to **GSS-SPNEGO**.

Additionally, if channel binding is enforced on the AD side, configure any systems that use SASL with SSL/TLS in the following way:

1. Install the latest versions of the **cyrus-sasl**, **openldap** and **krb5-libs** packages that are shipped with RHEL 8.3 and later.
2. In the **/etc/openldap/ldap.conf** file, specify the correct channel binding type by setting the **SASL_CBINDING** option to **tls-endpoint**.

For more information, see [Impact of Microsoft Security Advisory ADV190023 | LDAP Channel Binding and LDAP Signing on RHEL and AD integration](#).

(BZ#1873567)

SSSD, adcli, and realmd now support the deprecated RC4 cipher suite with a new system-wide cryptographic subpolicy

This update introduces the new **AD-SUPPORT** cryptographic subpolicy that enables the Rivest Cipher 4 (RC4) cipher suite for the following utilities:

- the System Security Services Daemon (SSSD)
- **adcli**
- **realmd**

As an administrator, you can activate the new **AD-SUPPORT** subpolicy when Active Directory (AD) is not configured to use Advanced Encryption Standard (AES) in the following scenarios:

- SSSD is used on a RHEL system connected directly to AD.
- **adcli** is used to join an AD domain or to update host attributes, for example the host key.
- **realmd** is used to join an AD domain.

Red Hat recommends enabling the new subpolicy if one of the following conditions applies:

- The user or service accounts in AD have RC4 encryption keys and lack AES encryption keys.

- The trust links between individual Active Directory domains have RC4 encryption keys and lack AES encryption keys.

To enable the **AD-SUPPORT** subpolicy in addition to the **DEFAULT** cryptographic policy, enter:

```
# update-crypto-policies --set DEFAULT:AD-SUPPORT
```

([BZ#1866695](#))

authselect has a new minimal profile

The **authselect** utility has a new **minimal** profile. You can use this profile to serve only local users and groups directly from system files instead of using other authentication providers. Therefore, you can safely remove the **SSSD**, **winbind**, and **fpprintd** packages and can use this profile on systems that require minimal installation to save disk and memory space.

([BZ#1654018](#))

SSSD now updates Samba's secrets.tdb file when rotating a password

A new **ad_update_samba_machine_account_password** option in the **sssd.conf** file is now available in RHEL. You can use it to set SSSD to automatically update the Samba **secrets.tdb** file when rotating a machine's domain password while using Samba.

However, if SELinux is in enforcing mode, SSSD fails to update the **secrets.tdb** file. Consequently, Samba does not have access to the new password. To work around this problem, set SELinux to permissive mode.

([BZ#1793727](#))

SSSD now enforces AD GPOs by default

The default setting for the SSSD option **ad_gpo_access_control** is now **enforcing**. In RHEL 8, SSSD enforces access control rules based on Active Directory Group Policy Objects (GPOs) by default.

Red Hat recommends ensuring GPOs are configured correctly in Active Directory before upgrading from RHEL 7 to RHEL 8. If you would not like to enforce GPOs, change the value of the **ad_gpo_access_control** option in the **/etc/sss/sss.conf** file to **permissive**.

([JIRA:RHELPLAN-51289](#))

Directory Server now supports the pwdReset operation attribute

This enhancement adds support for the **pwdReset** operation attribute to Directory Server. When an administrator changes the password of a user, Directory Server sets **pwdReset** in the user's entry to **true**. As a result, applications can use this attribute to identify if a password of a user has been reset by an administrator.

Note that **pwdReset** is an operational attribute and, therefore, users cannot edit it.

([BZ#1775285](#))

Directory Server now logs the work and operation time in RESULT entries

With this update, Directory Server now logs two additional time values in **RESULT`entries in the`** **/var/log/dirsrv/slaped-`<instance_name>`/access** file:

- The **wtime** value indicates how long it took for an operation to move from the work queue to a worker thread.
- The **optime** value shows the time the actual operation took to be completed once a worker thread started the operation.

The new values provide additional information about how the Directory Server handles load and processes operations.

For further details, see the [Access Log Reference](#) section in the Red Hat Directory Server Configuration, Command, and File Reference.

([BZ#1850275](#))

5.1.14. Desktop

Single-application session is now available

You can now start GNOME in a single-application session, also known as kiosk mode. In this session, GNOME displays only a full-screen window of an application that you have configured.

To enable the single-application session:

1. Install the **gnome-session-kiosk-session** package:

```
# yum install gnome-session-kiosk-session
```

2. Create and edit the **\$HOME/.local/bin/redhat-kiosk** file of the user that will open the single-application session.
In the file, enter the executable name of the application that you want to launch.

For example, to launch the **Text Editor** application:

```
#!/bin/sh  
  
gedit &
```

3. Make the file executable:

```
$ chmod +x $HOME/.local/bin/redhat-kiosk
```

4. At the GNOME login screen, select the **Kiosk** session from the cogwheel button menu and log in as the single-application user.

([BZ#1739556](#))

tigervnc has been rebased to version 1.10.1

The **tigervnc** suite has been rebased to version 1.10.1. The update contains number of fixes and improvements. Most notably:

- tigervnc now only supports starting of the virtual network computing (VNC) server using the **systemd** service manager.
- The clipboard now supports full Unicode in the native viewer, **WinVNC** and **Xvnc/libvnc.so**.

- The native client will now respect the system trust store when verifying server certificates.
- The Java web server has been removed.
- **x0vncserver** can now be configured to only allow local connections.
- **x0vncserver** has received fixes for when only part of the display is shared.
- Polling is now default in **WinVNC**.
- Compatibility with VMware's VNC server has been improved.
- Compatibility with some input methods on macOS has been improved.
- Automatic "repair" of JPEG artefacts has been improved.

([BZ#1806992](#))

5.1.15. Graphics infrastructures

Support for new graphics cards

The following graphics cards are now fully supported:

- The AMD Navi 14 family, which includes the following models:
 - Radeon RX 5300
 - Radeon RX 5300 XT
 - Radeon RX 5500
 - Radeon RX 5500 XT
- The AMD Renoir APU family, which includes the following models:
 - Ryzen 3 4300U
 - Ryzen 5 4500U, 4600U, and 4600H
 - Ryzen 7 4700U, 4800U, and 4800H
- The AMD Dali APU family, which includes the following models:
 - Athlon Silver 3050U
 - Athlon Gold 3150U
 - Ryzen 3 3250U

Additionally, the following graphics drivers have been updated:

- The Matrox **mgag200** driver

([JIRA:RHELPLAN-55009](#))

Hardware acceleration with Nvidia Volta and Turing

The **nouveau** graphics driver now supports hardware acceleration with the Nvidia Volta and Turing GPU families. As a result, the desktop and applications that use 3D graphics now render efficiently on the GPU. Additionally, this frees the CPU for other tasks and improves the overall system responsiveness.

(JIRA:RHELPLAN-57564)

Reduced display tearing on XWayland

The XWayland display back end now enables the XPresent extension. Using XPresent, applications can efficiently update their window content, which reduces display tearing.

This feature significantly improves the user interface rendering of full-screen OpenGL applications, such as 3D editors.

(JIRA:RHELPLAN-57567)

Intel Tiger Lake GPUs are now supported

This update adds support for the Intel Tiger Lake family of GPUs. This includes Intel UHD Graphics and Intel Xe GPUs found with the following CPU models:

<https://ark.intel.com/content/www/us/en/ark/products/codename/88759/tiger-lake.html>.

You no longer have to set the **i915.alpha_support=1** or **i915.force_probe=*** kernel option to enable Tiger Lake GPU support.

This enhancement was released as part of the [RHSA-2021:0558](#) asynchronous advisory.

(BZ#1882620)

5.1.16. The web console

Setting privileges from within the web console session

With this update the web console provides an option to switch between administrative access and limited access from inside of a user session. You can switch between the modes by clicking the **Administrative access** or **Limited access** indicator in your web console session.

(JIRA:RHELPLAN-42395)

Improvements to logs searching

With this update, the web console introduces a search box that supports several new ways of how the users can search among logs. The search box supports regular expression searching in log messages, specifying service or searching for entries with specific log fields.

([BZ#1710731](#))

Overview page shows more detailed Insights reports

With this update, when a machine is connected to Red Hat Insights, the **Health** card in the **Overview** page in the web console shows more detailed information about number of hits and their priority.

(JIRA:RHELPLAN-42396)

5.1.17. Red Hat Enterprise Linux System Roles

Terminal log role added to RHEL System Roles

With this enhancement, a new *Terminal log* (TLOG) role has been added to RHEL system roles shipped with the **rhel-system-roles** package. Users can now use the **tlog** role to setup and configure session recording using Ansible.

Currently, the **tlog** role supports the following tasks:

- Configure **tlog** to log recording data to the **systemd** journal
- Enable session recording for explicit users and groups, via SSSD

([BZ#1822158](#))

RHEL Logging System Role is now available for Ansible

With the Logging System Role, you can deploy various logging configurations consistently on local and remote hosts. You can configure a RHEL host as a server to collect logs from many client systems.

([BZ#1677739](#))

rhel-system-roles-sap fully supported

The **rhel-system-roles-sap** package, previously available as a Technology Preview, is now fully supported. It provides Red Hat Enterprise Linux (RHEL) System Roles for SAP, which can be used to automate the configuration of a RHEL system to run SAP workloads. These roles greatly reduce the time to configure a system to run SAP workloads by automatically applying the optimal settings that are based on best practices outlined in relevant SAP Notes. Access is limited to RHEL for SAP Solutions offerings. Please contact Red Hat Customer Support if you need assistance with your subscription.

The following new roles in the **rhel-system-roles-sap** package are fully supported:

- **sap-preconfigure**
- **sap-netweaver-preconfigure**
- **sap-hana-preconfigure**

For more information, see [Red Hat Enterprise Linux System Roles for SAP](#) .

([BZ#1660832](#))

The metrics RHEL System Role is now available for Ansible.

With the **metrics** RHEL System Role, you can configure, for local and remote hosts:

- performance analysis services via the **pcp** application
- visualisation of this data using a **grafana** server
- querying of this data using the **redis** data source without having to manually configure these services separately.

([BZ#1890499](#))

rhel-system-roles-sap upgraded

The **rhel-system-roles-sap** packages have been upgraded to upstream version 2.0.0, which provides multiple bug fixes and enhancements. Notable changes include:

- Improve hostname configuration and checking

- Improve **uuid** status detection and handling
- Add support for the **--check (-c)** option
- Increase **nofile** limits from 32800 to 65536
- Add the **nfs-utils** file to **sap_preconfigure_packages***
- Disable **firewalld**. With this change we disable **firewalld** only when it is installed.
- Add minimum required versions of the **setup** package for RHEL 8.0 and RHEL 8.1.
- Improve the **tmpfiles.d/sap.conf** file handling
- Support single step execution or checking of SAP notes
- Add the required **compat-sap-c++** packages
- Improve minimum package installation handling
- Detect if a reboot is required after applying the RHEL System Roles
- Support setting any SELinux state. Default state is **"disabled"**
- No longer fail if there is more than one line with identical IP addresses
- No longer modify **/etc/hosts** if there is more than one line containing **sap_ip**
- Support for HANA on RHEL 7.7
- Support for adding a repository for the IBM service and productivity tools for Power, required for SAP HANA on the **ppc64le** platform

([BZ#1844190](#))

The **storage** RHEL System Role now supports file system management

With this enhancement, administrators can use the **storage** RHEL System Role to:

- resize an **ext4** file
- resize a LVM file
- create a swap partition, if it does not exist, or to modify the swap partition, if it already exists, on a block device using the default parameters.

([BZ#1959289](#))

5.1.18. Virtualization

Migrating a virtual machine to a host with incompatible TSC setting now fails faster

Previously, migrating a virtual machine to a host with incompatible Time Stamp Counter (TSC) setting failed late in the process. With this update, attempting such a migration generates an error before the migration process starts.

([JIRA:RHELPLAN-45950](#))

Virtualization support for 2nd generation AMD EPYC processors

With this update, virtualization on RHEL 8 adds support for the 2nd generation AMD EPYC processors, also known as EPYC Rome. As a result, virtual machines hosted on RHEL 8 can now use the **EPYC-Rome** CPU model and utilise new features that the processors provide.

(JIRA:RHELPLAN-45959)

New command: `virsh iothreadset`

This update introduces the **virsh iothreadset** command, which can be used to configure dynamic IOThread polling. This makes it possible to set up virtual machines with lower latencies for I/O-intensive workloads at the expense of greater CPU consumption for the IOThread. For specific options, see the `virsh` man page.

(JIRA:RHELPLAN-45958)

UMIP is now supported by KVM on 10th generation Intel Core processors

With this update, the User-mode Instruction Prevention (UMIP) feature is now supported by KVM for hosts running on 10th generation Intel Core processors, also known as Ice Lake Servers. The UMIP feature issues a general protection exception if certain instructions, such as **sgdt**, **sidt**, **sldt**, **smsw**, and **str**, are executed when the Current Privilege Level (CPL) is greater than 0. As a result, UMIP ensures system security by preventing unauthorized applications from accessing certain system-wide settings which can be used to initiate privilege escalation attacks.

(JIRA:RHELPLAN-45957)

The libvirt library now supports Memory Bandwidth Allocation

libvirt now supports Memory Bandwidth Allocation (MBA). With MBA, you can allocate parts of host memory bandwidth in vCPU threads by using the **<memorytune>** element in the **<cputune>** section.

MBA is an extension of the existing Cache QoS Enforcement (CQE) feature found in the Intel Xeon v4 processors, also known as Broadwell server. For tasks that are associated with the CPU affinity, the mechanism used by MBA is the same as in CQE.

(JIRA:RHELPLAN-45956)

RHEL 6 virtual machines now support the Q35 machine type

Virtual machines (VMs) hosted on RHEL 8 that use RHEL 6 as their guest OS can now use Q35, a more modern PCI Express-based machine type. This provides a variety of improvements in features and performance of virtual devices, and ensures that a wider range of modern devices are compatible with RHEL 6 VMs.

(JIRA:RHELPLAN-45952)

All logged QEMU events now have a time stamp. As a result, users can more easily troubleshoot their virtual machines using logs saved in the `/var/log/libvirt/qemu/` directory.

QEMU logs now include time stamps for spice-server events

This update adds time stamps to ``spice-server`` event logs. Therefore, all logged QEMU events now have a time stamp. As a result, users can more easily troubleshoot their virtual machines using logs saved in the `/var/log/libvirt/qemu/` directory.

(JIRA:RHELPLAN-45945)

The bochs-display device is now supported

RHEL 8.3 and later introduce the Bochs display device, which is more secure than the currently used **stdvga** device. Note that all virtual machines (VMs) compatible with **bochs-display** will use it by default. This mainly includes VMs that use the UEFI interface.

(JIRA:RHELPLAN-45939)

Optimized MDS protection for virtual machines

With this update, a RHEL 8 host can inform its virtual machines (VMs) whether they are vulnerable to [Microarchitectural Data Sampling](#) (MDS). VMs that are not vulnerable do not use measures against MDS, which improves their performance.

(JIRA:RHELPLAN-45937)

Creating QCOW2 disk images on RBD now supported

With this update, it is possible to create QCOW2 disk images on RADOS Block Device (RBD) storage. As a result, virtual machines can use RBD servers for their storage back ends with QCOW2 images.

Note, however, that the write performance of QCOW2 disk images on RBD storage is currently lower than intended.

(JIRA:RHELPLAN-45936)

Maximum supported VFIO devices increased to 64

With this update, you can attach up to 64 PCI devices that use VFIO to a single virtual machine on a RHEL 8 host. This is up from 32 in RHEL 8.2 and prior.

(JIRA:RHELPLAN-45930)

discard and write-zeroes commands are now supported in QEMU/KVM

With this update, the **discard** and **write-zeroes** commands for **virtio-blk** are now supported in QEMU/KVM. As a result, virtual machines can use the **virtio-blk** device to discard unused sectors of an SSD, fill sectors with zeroes when they are emptied, or both. This can be used to increase SSD performance or to ensure that a drive is securely erased.

(JIRA:RHELPLAN-45926)

RHEL 8 now supports IBM POWER 9 XIVE

This update introduces support for the External Interrupt Virtualization Engine (XIVE) feature of IBM POWER9 to RHEL 8. As a result, virtual machines (VMs) running on a RHEL 8 hypervisor on an IBM POWER 9 system can use XIVE, which improves the performance of I/O-intensive VMs.

(JIRA:RHELPLAN-45922)

Control Group v2 support for virtual machines

With this update, the libvirt suite supports control groups v2. As a result, virtual machines hosted on RHEL 8 can take advantage of resource control capabilities of control group v2.

(JIRA:RHELPLAN-45920)

Paravirtualized IPIs are now supported for Windows virtual machines

With this update, the **hv_ipi** flag has been added to the supported hypervisor enlightenments for Windows virtual machines (VMs). This allows inter-processor interrupts (IPIs) to be sent via a hypercall. As a result, IPIs can be performed faster on VMs running a Windows OS.

(JIRA:RHELPLAN-45918)

Migrating virtual machines with enabled disk cache is now possible

This update makes the RHEL 8 KVM hypervisor compatible with disk cache live migration. As a result, it is now possible to live-migrate virtual machines with disk cache enabled.

(JIRA:RHELPLAN-45916)

macvtap interfaces can now be used by virtual machines in non-privileged sessions

It is now possible for virtual machines (VMs) to use a macvtap interface previously created by a privileged process. Notably, this enables VMs started by the non-privileged **user** session of **libvirt** to use a macvtap interface.

To do so, first create a macvtap interface in a privileged environment and set it to be owned by the user who will be running **libvirt** in a non-privileged session. You can do this using a management application such as the web console, or using command-line utilities as root, for example:

```
# ip link add link en2 name mymacvtap0 address 52:54:00:11:11:11 type macvtap mode bridge
# chown myuser /dev/tap$(cat /sys/class/net/mymacvtap0/ifindex)
# ip link set mymacvtap0 up
```

Afterwards, modify the **<target>** sub-element of the VM's **<interface>** configuration to reference the newly created macvtap interface:

```
<interface type='ethernet'>
  <model type='virtio' />
  <mac address='52:54:00:11:11:11' />
  <target dev='mymacvtap0' managed='no' />
</interface>
```

With this configuration, if **libvirt** is run as the user **myuser**, the VM will use the existing macvtap interface when started.

(JIRA:RHELPLAN-45915)

Virtual machines can now use features of 10th generation Intel Core processors

The **Icelake-Server** and **Icelake-Client** CPU model names are now available for virtual machines (VMs). On hosts with 10th generation Intel Core processors, using **Icelake-Server** or **Icelake-Client** as the CPU type in the XML configuration of a VM makes new features of these CPUs exposed to the VM.

(JIRA:RHELPLAN-45911)

QEMU now supports LUKS encryption

With this update, it is possible to create virtual disks using Linux Unified Key Setup (LUKS) encryption. You can encrypt the disks when creating the storage volume by including the **<encryption>** field in the virtual machine's (VM) XML configuration. You can also make the **LUKS** encrypted virtual disk completely transparent to the VM by including the **<encryption>** field in the disk's domain definition in the XML configuration file.

(JIRA:RHELPLAN-45910)

Improved logs for **nbdkit**

The **nbdkit** service logging has been modified to be less verbose. As a result, **nbdkit** logs only potentially important messages, and the logs created during **virt-v2v** conversions are shorter and easier to parse.

(JIRA:RHELPLAN-45909)

Improved consistency for virtual machines SELinux security labels and permissions

With this update, the **libvirt** service can record SELinux security labels and permissions associated with files, and restore the labels after modifying the files. As a result, for example, using **libguestfs** utilities to modify a virtual machine (VM) disk image owned by a specific user no longer changes the image owner to root.

Note that this feature does not work on file systems that do not support extended file attributes, such as NFS.

(JIRA:RHELPLAN-45908)

QEMU now uses the **gcrypt** library for XTS ciphers

With this update, the QEMU emulator has been changed to use the XTS cipher mode implementation provided by the **gcrypt** library. This improves the I/O performance of virtual machines whose host storage uses QEMU's native **luks** encryption driver.

(JIRA:RHELPLAN-45904)

Windows Virtio drivers can now be updated using Windows Updates

With this update, a new standard **SMBIOS** string is initiated by default when QEMU starts. The parameters provided in the **SMBIOS** fields make it possible to generate IDs for the virtual hardware running on the virtual machine(VM). As a result, Windows Update can identify the virtual hardware and the RHEL hypervisor machine type, and update the Virtio drivers on VMs running Windows 10+, Windows Server 2016, and Windows Server 2019+.

(JIRA:RHELPLAN-45901)

New command: **virsh guestinfo**

The **virsh guestinfo** command has been introduced to RHEL 8.3. This makes it possible to report the following types of information about a virtual machine (VM):

- Guest OS and file system information
- Active users
- The time zone used

Before running **virsh guestinfo**, ensure that the *qemu-guest-agent* package is installed. In addition, the **guest_agent** channel must be enabled in the VM's XML configuration, for example as follows:

```
<channel type='unix'>
  <target type='virtio' name='org.qemu.guest_agent.0'>
</channel>
```

(JIRA:RHELPLAN-45900)

VNNI for **BFLOAT16** inputs are now supported by KVM

With this update, Vector Neural Network Instructions (VNNI) supporting **BFLOAT16** inputs, also known as **AVX512_BF16** instructions, are now supported by KVM for hosts running on the 3rd Gen Intel Xeon scalable processors, also known as Cooper Lake. As a result, guest software can now use the **AVX512_BF16** instructions inside virtual machines, by enabling it in the virtual CPU configuration.

(JIRA:RHELPLAN-45899)

New command: **virsh pool-capabilities**

RHEL 8.3 introduces the **virsh pool-capabilities** command option. This command displays information that can be used for creating storage pools, as well as storage volumes within each pool, on your host. This includes:

- Storage pool types
- Storage pool source formats
- Target storage volume format types

(JIRA:RHELPLAN-45884)

Support for **CPUID.1F** in virtual machines with Intel Xeon Platinum 9200 series processors

With this update, virtual machines hosted on RHEL 8 can be configured with a virtual CPU topology of multiple dies, using the Extended Topology Enumeration leaf feature (CPUID.1F). This feature is supported by Intel Xeon Platinum 9200 series processors, previously known as Cascade Lake. As a result, it is now possible on hosts that use Intel Xeon Platinum 9200 series processors to create a vCPU topology that mirrors the physical CPU topology of the host.

(JIRA:RHELPLAN-37573, JIRA:RHELPLAN-45934)

Virtual machines can now use features of 3rd Generation Intel Xeon Scalable Processors

The **Cooperlake** CPU model name is now available for virtual machines (VMs). Using **Cooperlake** as the CPU type in the XML configuration of a VM makes new features from the 3rd Generation Intel Xeon Scalable Processors exposed to the VM, if the host uses this CPU.

(JIRA:RHELPLAN-37570)

Intel Optane persistent memory now supported by KVM

With this update, virtual machines hosted on RHEL 8 can benefit from the Intel Optane persistent memory technology, previously known as Intel Crystal Ridge. Intel Optane persistent memory storage devices provide data center-class persistent memory technology, which can significantly increase transaction throughput.

(JIRA:RHELPLAN-14068)

Virtual machines can now use Intel Processor Trace

With this update, virtual machines (VMs) hosted on RHEL 8 are able to use the Intel Processor Trace (PT) feature. When your host uses a CPU that supports Intel PT, you can use specialized Intel software to collect a variety of metrics about the performance of your VM's CPU. Note that this also requires enabling the **intel-pt** feature in the XML configuration of the VM.

(JIRA:RHELPLAN-7788)

DASD devices can now be assigned to virtual machines on IBM Z

Direct-access storage devices (DASDs) provide a number of specific storage features. Using the **vfio-ccw** feature, you can assign DASDs as mediated devices to your virtual machines (VMs) on IBM Z hosts. This for example makes it possible for the VM to access a z/OS dataset, or to share the assigned DASDs with a z/OS machine.

(JIRA:RHELPLAN-40234)

IBM Secure Execution supported for IBM Z

When using IBM Z hardware to run your RHEL 8 host, you can improve the security of your virtual machines (VMs) by configuring IBM Secure Execution for the VMs. IBM Secure Execution, also known as Protected Virtualization, prevents the host system from accessing a VM's state and memory contents.

As a result, even if the host is compromised, it cannot be used as a vector for attacking the guest operating system. In addition, Secure Execution can be used to prevent untrusted hosts from obtaining sensitive information from the VM.

(JIRA:RHELPLAN-14754)

5.1.19. RHEL in cloud environments

cloud-utils-growpart rebased to 0.31

The **cloud-utils-growpart** package has been upgraded to version 0.31, which provides multiple bug fixes and enhancements. Notable changes include:

- A bug that prevented GPT disks from being grown past 2TB has been fixed.
- The **growpart** operation no longer fails when the start sector and size are the same.
- Resizing a partition using the **sgdisk** utility previously in some cases failed. This problem has now been fixed.

([BZ#1846246](#))

5.1.20. Containers

skopeo container image is now available

The **registry.redhat.io/rhel8/skopeo** container image is a containerized implementation of the **skopeo** package. The **skopeo** tool is a command-line utility that performs various operations on container images and image repositories. This container image allows you to inspect container images in a registry, to remove a container image from a registry, and to copy container images from one unauthenticated container registry to another. To pull the **registry.redhat.io/rhel8/skopeo** container image, you need an active Red Hat Enterprise Linux subscription.

([BZ#1627900](#))

buildah container image is now available

The **registry.redhat.io/rhel8/buildah** container image is a containerized implementation of the **buildah** package. The **buildah** tool facilitates building OCI container images. This container image allows you to build container images without the need to install the **buildah** package on your system. The use-case does not cover running this image in rootless mode as a non-root user. To pull the **registry.redhat.io/rhel8/buildah** container image, you need an active Red Hat Enterprise Linux subscription.

([BZ#1627898](#))

Podman v2.0 RESTful API is now available

The new REST based Podman 2.0 API replaces the old remote API based on the varlink library. The new API works in both a rootful and a rootless environment and provides a docker compatibility layer.

(JIRA:RHELPLAN-37517)

Installing Podman does not require **container-selinux**

With this enhancement, the installation of the **container-selinux** package is now optional during the container build. As a result, Podman has fewer dependencies on other packages.

([BZ#1806044](#))

5.2. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS

This chapter provides system administrators with a summary of significant changes in the kernel shipped with Red Hat Enterprise Linux 8.3. These changes could include for example added or updated **proc** entries, **sysctl**, and **sysfs** default values, boot parameters, kernel configuration options, or any noticeable behavior changes.

New kernel parameters

acpi_no_watchdog = [HW,ACPI,WDT]

This parameter enables to ignore the Advanced Configuration and Power Interface (ACPI) based watchdog interface (WDAT) and let the native driver control the watchdog device instead.

dfltcc = [HW,S390]

This parameter configures the **zlib** hardware support for IBM Z architectures.

Format: { on | off | def_only | inf_only | always }

The options are:

- **on** (default) - IBM Z **zlib** hardware support for compression on level 1 and decompression
- **off** - No IBM Z **zlib** hardware support
- **def_only** - IBM Z **zlib** hardware support for the **deflate** algorithm only (compression on level 1)
- **inf_only** - IBM Z **zlib** hardware support for the **inflate** algorithm only (decompression)
- **always** - Similar as **on**, but ignores the selected compression level and always uses hardware support (used for debugging)

irqchip.gicv3_pseudo_nmi = [ARM64]

This parameter enables support for pseudo non-maskable interrupts (NMIs) in the kernel.

To use this parameter you need to build the kernel with the **CONFIG_ARM64_PSEUDO_NMI** configuration item.

panic_on_taint =

Bitmask for conditionally calling **panic()** in **add_taint()**

Format: <hex>[,nouser_taint]

A hexadecimal bitmask which represents a set of **TAINT** flags that will cause the kernel to panic when the **add_taint()** system call is invoked with any of the flags in this set. The optional **nousertaint** switch prevents userspace-forced crashes by writing to the **/proc/sys/kernel/tainted** file any flagset that matches the bitmask in **panic_on_taint**.

For more information see the [upstream documentation](#).

prot_virt = [S390]

Format: <bool>

This parameter enables hosting of protected virtual machines which are isolated from the hypervisor if the hardware support is present.

rcutree.use_softirq = [KNL]

This parameter enables elimination of Tree-RCU **softirq** processing.

If you set this parameter to zero, it moves all **RCU_SOFTIRQ** processing to per-CPU rcuc kthreads.

If you set **rcutree.use_softirq** to a non-zero value (default), **RCU_SOFTIRQ** is used by default.

Specify **rcutree.use_softirq=0** to use rcuc kthreads.

split_lock_detect = [X86]

This parameter enables the split lock detection. When enabled, and if hardware support is present, atomic instructions that access data across cache line boundaries will result in an alignment check exception.

The options are:

- **off** - not enabled
- **warn** - the kernel will emit rate limited warnings about applications that trigger the Alignment Check Exception (#AC). This mode is the default on CPUs that supports split lock detection.
- **fatal** - the kernel will send Bus error (SIGBUS) signal to applications that trigger the #AC exception.
If the #AC exception is hit while not executing in the user mode, the kernel will issue an oops error in either the **warn** or **fatal** mode.

srbds = [X86,INTEL]

This parameter controls the Special Register Buffer Data Sampling (SRBDS) mitigation.

Certain CPUs are vulnerable to a Microarchitectural Data Sampling (MDS)-like exploit which can leak bits from the random number generator.

By default, microcode mitigates this issue. However, the microcode fix can cause the **RDRAND** and **RDSEED** instructions to become much slower. Among other effects, this will result in reduced throughput from the **urandom** kernel random number source device.

To disable the microcode mitigation, set the following option:

- **off** - Disable mitigation and remove performance impact to **RDRAND** and **RDSEED**

svm = [PPC]

Format: { on | off | y | n | 1 | 0 }

This parameter controls the use of the Protected Execution Facility on pSeries systems.

nopv = [X86,XEN,KVM,HYPER_V,VMWARE]

This parameter disables the PV optimizations which forces the guest to run as generic guest with no PV drivers.

Currently supported are XEN HVM, KVM, HYPER_V and VMWARE guests.

Updated kernel parameters

hugepagesz = [HW]

This parameter specifies a huge page size. Use this parameter in conjunction with the **hugepages** parameter to pre-allocate a number of huge pages of the specified size.

Specify the **hugepagesz** and **hugepages** parameters in pairs such as:

```
hugepagesz=2M hugepages=512
```

The **hugepagesz** parameter can only be specified once on the command line for a specific huge page size. Valid huge page sizes are architecture dependent.

hugepages = [HW]

This parameter specifies the number of huge pages to pre-allocate. This parameter typically follows the valid **hugepagesz** or **default_hugepagesz** parameter.

However, if **hugepages** is the first or the only HugeTLB command-line parameter, it implicitly specifies the number of huge pages of the default size to allocate. If the number of huge pages of the default size is implicitly specified, it can not be overwritten by the **hugepagesz** + **hugepages** parameter pair for the default size.

For example, on an architecture with 2M default huge page size:

```
hugepages=256 hugepagesz=2M hugepages=512
```

Settings from the example above results in allocation of 256 2M huge pages and a warning message that the **hugepages=512** parameter was ignored. If **hugepages** is preceded by invalid **hugepagesz**, **hugepages** will be ignored.

default_hugepagesz = [HW]

This parameter specifies the default huge page size. You can specify **default_hugepagesz** only once on the command-line. Optionally, you can follow **default_hugepagesz** with the **hugepages** parameter to pre-allocate a specific number of huge pages of the default size. Also, you can implicitly specify the number of default-sized huge pages to pre-allocate.

For example, on an architecture with 2M default huge page size:

```
hugepages=256
default_hugepagesz=2M hugepages=256
hugepages=256 default_hugepagesz=2M
```

Settings from the example above all results in allocation of 256 2M huge pages. Valid default huge page size is architecture dependent.

efi = [EFI]

Format: { "old_map", "nochunk", "noruntime", "debug", "nosoftreserve" }

The options are:

- **old_map** [X86-64] - Switch to the old ioremap-based EFI runtime services mapping. 32-bit still uses this one by default

- **nochunk** - Disable reading files in "chunks" in the EFI boot stub, as chunking can cause problems with some firmware implementations
- **noruntime** - Disable EFI runtime services support
- **debug** - Enable miscellaneous debug output
- **nosoftreserve** - The **EFI_MEMORY_SP** (Specific Purpose) attribute sometimes causes the kernel to reserve the memory range for a memory mapping driver to claim. Specify **efi=nosoftreserve** to disable this reservation and treat the memory by its base type (for example **EFI_CONVENTIONAL_MEMORY** / "System RAM").

intel_iommu = [DMAR]

Intel IOMMU driver Direct Memory Access Remapping (DMAR).

The added options are:

- **nobounce** (Default off) - Disable bounce buffer for untrusted devices such as the Thunderbolt devices. This will treat the untrusted devices as the trusted ones. Hence this setting might expose security risks of direct memory access (DMA) attacks.

mem = nn[KMG] [KNL,BOOT]

This parameter forces the usage of a specific amount of memory.

The amount of memory to be used in cases as follows:

1. For test.
2. When the kernel is not able to see the whole system memory.
3. Memory that lies after the **mem** boundary is excluded from the hypervisor, then assigned to KVM guests.
[X86] Work as limiting max address. Use together with the **memmap** parameter to avoid physical address space collisions. Without **memmap**, Peripheral Component Interconnect (PCI) devices could be placed at addresses belonging to unused RAM.

Note that this setting only takes effect during the boot time since in the case 3 above, the memory may need to be hot added after the boot if the system memory of hypervisor is not sufficient.

pci = [PCI]

Various Peripheral Component Interconnect (PCI) subsystem options.

Some options herein operate on a specific device or a set of devices (**<pci_dev>**). These are specified in one of the following formats:

```
[<domain>:]<bus>:<dev>.<func>[/<dev>.<func>]*  
pci:<vendor>:<device>[:<subvendor>:<subdevice>]
```

Note that the first format specifies a PCI bus/device/function address which may change if new hardware is inserted, if motherboard firmware changes, or due to changes caused by other kernel parameters. If the domain is left unspecified, it is taken to be zero. Optionally, a path to a device through multiple device/function addresses can be specified after the base address (this is more robust against renumbering issues). The second format selects devices using IDs from the configuration space which may match multiple devices in the system.

The options are:

- **hpmmiosize** - The fixed amount of bus space which is reserved for hotplug bridge's Memory-mapped I/O (MMIO) window. The default size is 2 megabytes.
- **hpmmioprefsize** - The fixed amount of bus space which is reserved for hotplug bridge's MMIO_PREF window. The default size is 2 megabytes.

pcie_ports = [PCIE]

Peripheral Component Interconnect Express (PCIe) port services handling.

The options are:

- **native** - Use native PCIe services (PME, AER, DPC, PCIe hotplug) even if the platform does not give the OS permission to use them. This setting may cause conflicts if the platform also tries to use these services.
- **dpc-native** - Use native PCIe service for DPC only. This setting may cause conflicts if firmware uses AER or DPC.
- **compat** - Disable native PCIe services (PME, AER, DPC, PCIe hotplug).

rcu_nocbs = [KNL]

The argument is a CPU list. The string "all" can be used to specify every CPU on the system.

usbcore.authorized_default = [USB]

The default USB device authorization.

The options are:

- **-1** (Default) - Authorized except for wireless USB
- **0** - Not authorized
- **1** - Authorized
- **2** - Authorized if the device is connected to the internal port

usbcore.old_scheme_first = [USB]

This parameter enables to start with the old device initialization scheme. This setting applies only to low and full-speed devices (default 0 = off).

usbcore.quirks = [USB]

A list of quirk entries to augment the built-in USB core quirk list. The list entries are separated by commas. Each entry has the form **VendorID:ProductID:Flags**, for example

quirks=0781:5580:bk,0a5c:5834:gij. The IDs are 4-digit hex numbers and **Flags** is a set of letters. Each letter will change the built-in quirk; setting it if it is clear and clearing it if it is set.

The added flags:

- **o** - **USB_QUIRK_HUB_SLOW_RESET**, hub needs extra delay after resetting its port

New /proc/sys/fs parameters

protected_fifos

This parameter is based on the restrictions in the Openwall software and provides protection by allowing to avoid unintentional writes to an attacker-controlled FIFO where a program intended to create a regular file.

The options are:

- **0** - Writing to FIFOs is unrestricted.
- **1** - Does not allow the **O_CREAT** flag open on FIFOs that we do not own in world writable sticky directories unless they are owned by the owner of the directory.
- **2** - Applies to group writable sticky directories.

protected_regular

This parameter is similar to the **protected_fifos** parameter, however it avoids writes to an attacker-controlled regular file where a program intended to create one.

The options are:

- **0** - Writing to regular files is unrestricted.
- **1** - Does not allow the **O_CREAT** flag open on regular files that we do not own in world writable sticky directories unless they are owned by the owner of the directory.
- **2** - Applies to group writable sticky directories.

5.3. DEVICE DRIVERS

5.3.1. New drivers

Network drivers

- CAN driver for Kvaser CAN/USB devices (kvaser_usb.ko.xz)
- Driver for Theobroma Systems UCAN devices (ucan.ko.xz)
- Pensando Ethernet NIC Driver (ionic.ko.xz)

Graphics drivers and miscellaneous drivers

- Generic Remote Processor Framework (remoteproc.ko.xz)
- Package Level C-state Idle Injection for Intel® CPUs (intel_powerclamp.ko.xz)
- X86 PKG TEMP Thermal Driver (x86_pkg_temp_thermal.ko.xz)
- INT3402 Thermal driver (int3402_thermal.ko.xz)
- ACPI INT3403 thermal driver (int3403_thermal.ko.xz)
- Intel® acpi thermal rel misc dev driver (acpi_thermal_rel.ko.xz)
- INT3400 Thermal driver (int3400_thermal.ko.xz)
- Intel® INT340x common thermal zone handler (int340x_thermal_zone.ko.xz)
- Processor Thermal Reporting Device Driver (processor_thermal_device.ko.xz)
- Intel® PCH Thermal driver (intel_pch_thermal.ko.xz)
- DRM gem ttm helpers (drm_ttm_helper.ko.xz)

- Device node registration for cec drivers (cec.ko.xz)
- Fairchild FUSB302 Type-C Chip Driver (fusb302.ko.xz)
- VHOST IOTLB (vhost_iotlb.ko.xz)
- vDPA-based vhost backend for virtio (vhost_vdpa.ko.xz)
- VMware virtual PTP clock driver (ptp_vmw.ko.xz)
- Intel® LPSS PCI driver (intel-lpss-pci.ko.xz)
- Intel® LPSS core driver (intel-lpss.ko.xz)
- Intel® LPSS ACPI driver (intel-lpss-acpi.ko.xz)
- Mellanox watchdog driver (mlx_wdt.ko.xz)
- Mellanox FAN driver (mlxreg-fan.ko.xz)
- Mellanox regmap I/O access driver (mlxreg-io.ko.xz)
- Intel® speed select interface pci mailbox driver (isst_if_mbox_pci.ko.xz)
- Intel® speed select interface mailbox driver (isst_if_mbox_msr.ko.xz)
- Intel® speed select interface mmio driver (isst_if_mmio.ko.xz)
- Mellanox LED regmap driver (leds-mlxreg.ko.xz)
- vDPA Device Simulator (vdpa_sim.ko.xz)
- Intel® Tiger Lake PCH pinctrl/GPIO driver (pinctrl-tigerlake.ko.xz)
- PXA2xx SSP SPI Controller (spi-pxa2xx-platform.ko.xz)
- CE4100/LPSS PCI-SPI glue code for PXA's driver (spi-pxa2xx-pci.ko.xz)
- Hyper-V PCI Interface (pci-hyperv-intf.ko.xz)
- vDPA bus driver for virtio devices (virtio_vdpa.ko.xz)

5.3.2. Updated drivers

Network driver updates

- VMware vmxnet3 virtual NIC driver (vmxnet3.ko.xz) has been updated to version 1.5.0.0-k.
- Realtek RTL8152/RTL8153 Based USB Ethernet Adapters (r8152.ko.xz) has been updated to version 1.09.10.
- Broadcom BCM573xx network driver (bnxt_en.ko.xz) has been updated to version 1.10.1.
- The Netronome Flow Processor (NFP) driver (nfp.ko.xz) has been updated to version 4.18.0-240.el8.x86_64.
- Intel® Ethernet Switch Host Interface Driver (fm10k.ko.xz) has been updated to version 0.27.1-k.

- Intel® Ethernet Connection E800 Series Linux Driver (ice.ko.xz) has been updated to version 0.8.2-k.

Storage driver updates

- Emulex LightPulse Fibre Channel SCSI driver (lpfc.ko.xz) has been updated to version 0:12.8.0.1.
- QLogic FCoE Driver (bnx2fc.ko.xz) has been updated to version 2.12.13.
- LSI MPT Fusion SAS 3.0 Device Driver (mpt3sas.ko.xz) has been updated to version 34.100.00.00.
- Driver for HP Smart Array Controller version (hpsa.ko.xz) has been updated to version 3.4.20-170-RH5.
- QLogic Fibre Channel HBA Driver (qla2xxx.ko.xz) has been updated to version 10.01.00.25.08.3-k.
- Broadcom MegaRAID SAS Driver (megaraid_sas.ko.xz) has been updated to version 07.714.04.00-rh1.

Graphics and miscellaneous driver updates

- Standalone drm driver for the VMware SVGA device (vmwgfx.ko.xz) has been updated to version 2.17.0.0.
- Crypto Co-processor for Chelsio Terminator cards. (chcr.ko.xz) has been updated to version 1.0.0.0-ko.

5.4. BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 8.3 that have a significant impact on users.

5.4.1. Installer and image creation

RHEL 8 initial setup now works properly via SSH

Previously, the RHEL 8 initial setup interface did not display when logged in to the system using SSH. As a consequence, it was impossible to perform the initial setup on a RHEL 8 machine managed via SSH. This problem has been fixed, and RHEL 8 initial setup now works correctly when performed via SSH.

([BZ#1676439](#))

Installation failed when using the **reboot --kexec** command

Previously, the RHEL 8 installation failed when a Kickstart file that contained the **reboot --kexec** command was used.

With this update, the installation with **reboot --kexec** now works as expected.

([BZ#1672405](#))

America/New York time zone can now be set correctly

Previously, the interactive Anaconda installation process did not allow users to set the *America/New York* time zone when using a kickstart file. With this update, users can now set *America/New York* as the preferred time zone in the interactive installer if a time zone is not specified in the kickstart file.

(BZ#1665428)

SELinux contexts are now set correctly

Previously, when SELinux was in enforcing mode, incorrect SELinux contexts on some folders and files resulted in unexpected AVC denials when attempting to access these files after installation.

With this update, Anaconda sets the correct SELinux contexts. As a result, you can now access the folders and files without manually relabeling the filesystem.

(BZ#1775975)

Automatic partitioning now creates a valid **/boot** partition

Previously, when installing RHEL on a system using automatic partitioning or using a kickstart file with preconfigured partitions, the installer created a partitioning scheme that could contain an invalid **/boot** partition. Consequently, the automatic installation process ended prematurely because the verification of the partitioning scheme failed. With this update, Anaconda creates a partitioning scheme that contains a valid **/boot** partition. As a result, the automatic installation completes as expected.

(BZ#1630299)

A GUI installation using the Binary DVD ISO image now completes successfully without CDN registration

Previously, when performing a GUI installation using the Binary DVD ISO image file, a race condition in the installer prevented the installation from proceeding until you registered the system using the Connect to Red Hat feature.

With this update, you can now proceed with the installation without registering the system using the Connect to Red Hat feature.

(BZ#1823578)

iSCSI or FCoE devices created in Kickstart and used in **ignoredisk --only-use** command no longer stop the installation process

Previously, when the iSCSI or FCoE devices created in Kickstart were used in the **ignoredisk --only-use** command, the installation program failed with an error similar to **Disk "disk/by-id/scsi-360a9800042566643352b476d674a774a" given in ignoredisk command does not exist**. This stopped the installation process.

With this update, the problem has been fixed. The installation program continues working.

(BZ#1644662)

System registration using CDN failed with the error message **Name or service not known**

When you attempted to register a system using the Content Delivery Network (CDN), the registration process failed with the error message **Name or service not known**.

This issue occurred because the empty **Custom server URL** and **Custom Base URL** values overwrote the default values for system registration.

With this update, the empty values now do not overwrite the default values, and the system registration completes successfully.

(BZ#1862116)

5.4.2. Software management

dnf-automatic now updates only packages with correct GPG signatures

Previously, the **dnf-automatic** configuration file did not check GPG signatures of downloaded packages before performing an update. As a consequence, unsigned updates or updates signed by key which was not imported could be installed by **dnf-automatic** even though repository configuration requires GPG signature check (**gpgcheck=1**). With this update, the problem has been fixed, and **dnf-automatic** checks GPG signatures of downloaded packages before performing the update. As a result, only updates with correct GPG signatures are installed from repositories that require GPG signature check.

([BZ#1793298](#))

Trailing comma no longer causes entries removal in an `append` type option

Previously, adding a trailing comma (an empty entry at the end of the list) to an **append** type option (for example, **exclude**, **excludepkgs**, **includepkgs**) caused all entries in the option to be removed. Also, adding two commas (an empty entry) caused that only entries after the commas were used.

With this update, empty entries other than leading commas (an empty entry at the beginning of the list) are ignored. As a result, only the leading comma now removes existing entries from the **append** type option, and the user can use it to overwrite these entries.

([BZ#1788154](#))

5.4.3. Shells and command-line tools

The **ReaR disk layout no longer includes entries for Rancher 2 Longhorn iSCSI devices and file systems**

This update removes entries for Rancher 2 Longhorn iSCSI devices and file systems from the disk layout created by **ReaR**.

([BZ#1843809](#))

Rescue image creation with a file larger than 4 GB is now enabled on IBM POWER, little endian

Previously, the **ReaR** utility could not create rescue images containing files larger than 4GB on IBM POWER, little endian architecture. With this update, the problem has been fixed, and it is now possible to create a rescue image with a file larger than 4 GB on IBM POWER, little endian.

([BZ#1729502](#))

5.4.4. Security

SELinux no longer prevents **systemd-journal-gatewayd to call **newfstatat()** on `/dev/shm/` files used by **corosync****

Previously, SELinux policy did not contain a rule that allows the **systemd-journal-gatewayd** daemon to access files created by the **corosync** service. As a consequence, SELinux denied **systemd-journal-gatewayd** to call the **newfstatat()** function on shared memory files created by **corosync**. With this update, SELinux no longer prevents **systemd-journal-gatewayd** to call **newfstatat()** on shared memory files created by **corosync**.

([BZ#1746398](#))

Libreswan now works with **seccomp=enabled** on all configurations

Prior to this update, the set of allowed syscalls in the **Libreswan** SECCOMP support implementation did not match new usage of RHEL libraries. Consequently, when SECCOMP was enabled in the **ipsec.conf** file, the syscall filtering rejected even syscalls required for the proper functioning of the **pluto** daemon; the daemon was killed, and the **ipsec** service was restarted. With this update, all newly required syscalls have been allowed, and **Libreswan** now works with the **seccomp=enabled** option correctly.

([BZ#1544463](#))

SELinux no longer prevents **auditd** to halt or power off the system

Previously, the SELinux policy did not contain a rule that allows the Audit daemon to start a **power_unit_file_t systemd** unit. Consequently, **auditd** could not halt or power off the system even when configured to do so in cases such as no space left on a logging disk partition.

This update of the **selinux-policy** packages adds the missing rule, and **auditd** can now properly halt and power off the system only with SELinux in enforcing mode.

([BZ#1826788](#))

IPTABLES_SAVE_ON_STOP now works correctly

Previously, the **IPTABLES_SAVE_ON_STOP** feature of the **iptables** service did not work because files with saved IP tables content received incorrect SELinux context. This prevented the **iptables** script from changing permissions, and the script subsequently failed to save the changes. This update defines a proper context for the **iptables.save** and **ip6tables.save** files, and creates a filename transition rule. As a consequence, the **IPTABLES_SAVE_ON_STOP** feature of the **iptables** service works correctly.

([BZ#1776873](#))

NSCD databases can now use different modes

Domains in the **nsswitch_domain** attribute are allowed access to Name Service Cache Daemon (NSCD) services. Each NSCD database is configured in the **nscd.conf** file, and the **shared** property determines whether the database uses Shared memory or Socket mode. Previously, all NSCD databases had to use the same access mode, depending on the **nscd_use_shm** boolean value. Now, using Unix stream socket is always allowed, and therefore different NSCD databases can use different modes.

([BZ#1772852](#))

The **oscap-ssh** utility now works correctly when scanning a remote system with **--sudo**

When performing a Security Content Automation Protocol (SCAP) scan of a remote system using the **oscap-ssh** tool with the **--sudo** option, the **oscap** tool on the remote system saves scan result files and report files into a temporary directory as the **root** user. Previously, if the **umask** settings on the remote machine were changed, **oscap-ssh** might have been prevented access to these files. This update fixes the issue, and as a result, **oscap** saves the files as the target user, and **oscap-ssh** accesses the files normally.

([BZ#1803116](#))

OpenSCAP now handles remote file systems correctly

Previously, OpenSCAP did not reliably detect remote file systems if their mount specification did not start with two slashes. As a consequence, OpenSCAP handled some network-based file systems as local. With this update, OpenSCAP identifies file systems using the file-system type instead of the mount

specification. As a result, OpenSCAP now handles remote file systems correctly.

([BZ#1870087](#))

OpenSCAP no longer removes blank lines from YAML multi-line strings

Previously, OpenSCAP removed blank lines from YAML multi-line strings within generated Ansible remediations from a datastream. This affected Ansible remediations and caused the **openscap** utility to fail the corresponding Open Vulnerability and Assessment Language (OVAL) checks, producing false positive results. The issue is now fixed and as a result, **openscap** no longer removes blank lines from YAML multi-line strings.

([BZ#1795563](#))

OpenSCAP can now scan systems with large numbers of files without running out of memory

Previously, when scanning systems with low RAM and large numbers of files, the OpenSCAP scanner sometimes caused the system to run out of memory. With this update, OpenSCAP scanner memory management has been improved. As a result, the scanner no longer runs out of memory on systems with low RAM when scanning large numbers of files, for example package groups **Server with GUI** and **Workstation**.

([BZ#1824152](#))

config.enabled now controls statements correctly

Previously, the **rsyslog** incorrectly evaluated the **config.enabled** directive during the configuration processing of a statement. As a consequence, the **parameter not known** errors were displayed for each statement except for the **include()** one. With this update, the configuration is processed for all statements equally. As a result, **config.enabled** now correctly disables or enables statements without displaying any error.

([BZ#1659383](#))

fapolicyd no longer prevents RHEL updates

When an update replaces the binary of a running application, the kernel modifies the application binary path in memory by appending the " (deleted)" suffix. Previously, the **fapolicyd** file access policy daemon treated such applications as untrusted, and prevented them from opening and executing any other files. As a consequence, the system was sometimes unable to boot after applying updates.

With the release of the [RHBA-2020:5242](#) advisory, **fapolicyd** ignores the suffix in the binary path so the binary can match the trust database. As a result, **fapolicyd** enforces the rules correctly and the update process can finish.

([BZ#1897090](#))

The e8 profile can now be used to remediate RHEL 8 systems with **Server with GUI**

Using the OpenSCAP Anaconda Add-on to harden the system on the **Server With GUI** package group with profiles that select rules from the **Verify Integrity with RPM** group no longer requires an extreme amount of RAM on the system. The cause of this problem was the OpenSCAP scanner. For more details, see [Scanning large numbers of files with OpenSCAP causes systems to run out of memory](#) . As a consequence, the hardening of the system using the RHEL 8 Essential Eight (e8) profile now works also with **Server With GUI**.

([BZ#1816199](#))

5.4.5. Networking

Automatic loading of **iptables** extension modules by the **nft_compat** module no longer hangs

Previously, when the **nft_compat** module loaded an extension module while an operation on network name spaces (**netns**) happened in parallel, a lock collision could occur if that extension registered a **pernet** subsystem during initialization. As a consequence, the kernel-called **modprobe** command hang. This could also be caused by other services, such as **libvirtd**, that also execute **iptables** commands. This problem has been fixed. As a result, loading **iptables** extension modules by the **nft_compat** module no longer hangs.

(BZ#1757933)

The **firewalld** service now removes **ipsets** when the service stops

Previously, stopping the **firewalld** service did not remove **ipsets**. This update fixes the problem. As a result, **ipsets** are no longer left in the system after **firewalld** stops.

(BZ#1790948)

firewalld no longer retains **ipset** entries after shutdown

Previously, shutting down **firewalld** did not remove **ipset** entries. Consequently, **ipset** entries remained active in the kernel even after stopping the **firewalld** service. With this fix, shutting down **firewalld** removes **ipset** entries as expected.

(BZ#1682913)

firewalld now restores **ipset** entries after reloading

Previously, **firewalld** did not retain runtime **ipset** entries after reloading. Consequently, users had to manually add the missing entries again. With this update, **firewalld** has been modified to restore **ipset** entries after reloading.

(BZ#1809225)

nftables and **firewalld** services are now mutually exclusive

Previously, it was possible to enable **nftables** and **firewalld** services at the same time. As a consequence, **nftables** was overriding **firewalld** rulesets. With this update, **nftables** and **firewalld** services are now mutually exclusive so that these cannot be enabled at the same time.

(BZ#1817205)

5.4.6. Kernel

The **huge_page_setup_helper.py** script now works correctly

A patch that updated the **huge_page_setup_helper.py** script for Python 3 was accidentally removed. Consequently, after executing **huge_page_setup_helper.py**, the following error message appeared:

```
SyntaxError: Missing parentheses in call to 'print'
```

With this update, the problem has been fixed by updating the **libhugetlbfs.spec** file. As a result, **huge_page_setup_helper.py** does not show any error in the described scenario.

(BZ#1823398)

Systems with a large amount of persistent memory boot more quickly and without timeouts

Systems with a large amount of persistent memory took a long time to boot because the original source code allowed for just one initialization thread per node. For example, for a 4-node system there were 4 memory initialization threads. Consequently, if there were persistent memory file systems listed in the **/etc/fstab** file, the system could time out while waiting for devices to become available. With this update, the problem has been fixed because the source code now allows for multiple memory initialization threads within a single node. As a result, the systems boot more quickly and no timeouts appear in the described scenario.

(BZ#1666538)

The **bcc** scripts now successfully compile a BPF module

During the script code compilation to create a Berkeley Packet Filter (BPF) module, the **bcc** toolkit used kernel headers for data type definition. Some kernel headers needed the **KBUILD_MODNAME** macro to be defined. Consequently, those **bcc** scripts that did not add **KBUILD_MODNAME**, were likely to fail to compile a BPF module across various CPU architectures. The following **bcc** scripts were affected:

- **bindsnoop**
- **sofdsnoop**
- **solisten**
- **tcpaccept**
- **tcpconnect**
- **tcpconnlrat**
- **tcpdrop**
- **tcpretrans**
- **tcpsubnet**
- **tcptop**
- **tcptracer**

With this update, the problem has been fixed by adding **KBUILD_MODNAME** to the default **cflags** parameter for **bcc**. As a result, this problem no longer appears in the described scenario. Also, customer scripts do not need to define **KBUILD_MODNAME** themselves either.

(BZ#1837906)

bcc-tools and **bpfftrace** work properly on IBM Z

Previously, a feature backport introduced the **ARCH_HAS_NON_OVERLAPPING_ADDRESS_SPACE** kernel option. However, the **bcc-tools** package and **bpfftrace** tracing language package for IBM Z architectures did not have proper support for this option. Consequently, the **bpf()** system call failed with the **Invalid argument** exception and **bpfftrace** failed with an error stating **Error loading program** when trying to load the BPF program. With this update, the **ARCH_HAS_NON_OVERLAPPING_ADDRESS_SPACE** option is now removed. As a result, the problem no longer appears in the described scenario.

(BZ#1847837, BZ#1853964)

Boot process no longer fails due to lack of entropy

Previously, the boot process failed due to lack of entropy. A better mechanism is now used to allow the kernel to gather entropy early in the boot process, which does not depend on any hardware specific interrupts. This update fixes the problem by ensuring availability of sufficient entropy to secure random generation in early boot. As a result, the fix prevents kickstart timeout or slow boots and the boot process works as expected.

(BZ#1778762)

Repeated reboots using **kexec** now work as expected

Previously, during the kernel reboot on the Amazon EC2 Nitro platform, the remove module (**rmmod**) was not called during the **shutdown()** call of the kernel execution path. Consequently, repeated kernel reboots using the **kexec** system call led to a failure. With this update, the issue has been fixed by adding the PCI **shutdown()** handler that allows safe kernel execution. As a result, repeated reboots using **kexec** on Amazon EC2 Nitro platforms no longer fail.

(BZ#1758323)

Repeated reboots using vPMEM memory as dump target now works as expected

Previously, using Virtual Persistent Memory (vPMEM) namespaces as dump target for **kdump** or **fadump** caused the **papr_scm** module to unmap and remap the memory backed by vPMEM and re-add the memory to its linear map.

Consequently, this behavior triggered Hypervisor Calls (HCalls) to POWER Hypervisor. As a result, this slows down the capture kernel boot considerably and takes a long time to save the dump file. This update fixes the problem and the boot process now works as expected in the described scenario

(BZ#1792125)

Attempting to add ICE driver NIC port to a mode 5 bonding master interface no longer fails

Previously, attempting to add the **ICE** driver NIC port to a mode 5 (**balance-tlb**) bonding master interface led to a failure with an error **Master 'bond0', Slave 'ens1f0': Error: Enslave failed**. Consequently, you experienced an intermittent failure to add the NIC port to the bonding master interface. This update fixes the issue and adding the interface no longer fails.

(BZ#1791664)

The **cxgb4** driver no longer causes crash in the **kdump** kernel

Previously, the **kdump** kernel would crash while trying to save information in the **vmcore** file. Consequently, the **cxgb4** driver prevented the **kdump** kernel from saving a core for later analysis. To work around this problem, add the **novmcoredd** parameter to the **kdump** kernel command line to allow saving core files.

With the release of the [RHSA-2020:1769](#) advisory, the **kdump** kernel handles this situation properly and no longer crashes.

(BZ#1708456)

5.4.7. High availability and clusters

When a GFS2 file system is used with the Filesystem agent the **fast_stop** option now defaults to **no**

Previously, when a GFS2 file system was used with the Filesystem agent, the **fast_stop** option defaulted to **yes**. This value could result in unnecessary fence events due to the length of time it can take a GFS2 file system to unmount. With this update, this option defaults to **no**. For all other file systems it continues to default to **yes**.

(BZ#1814896)

fence_compute and **fence_evacuate** agents now interpret **insecure** option in a more standard way

Previously, the **fence_compute** and **fence_evacuate** agents worked as if **--insecure** was specified by default. With this update, customers who do not use valid certificates for their compute or evacuate services must set **insecure=true** and use the **--insecure** option when running manually from the CLI. This is consistent with the behavior of all other agents.

(BZ#1830776)

5.4.8. Dynamic programming languages, web and database servers

Optimized CPU consumption by **libdb**

A previous update to the **libdb** database caused an excessive CPU consumption in the trickle thread. With this update, the CPU usage has been optimized.

(BZ#1670768)

The **did_you_mean** Ruby gem no longer contains a file with a non-commercial license

Previously, the **did_you_mean** gem available in the **ruby:2.5** module stream contained a file with a non-commercial license. This update removes the affected file.

(BZ#1846113)

nginx can now load server certificates from hardware security tokens through the PKCS#11 URI

The **ssl_certificate** directive of the **nginx** web server supports loading TLS server certificates from hardware security tokens directly from PKCS#11 modules. Previously, it was impossible to load server certificates from hardware security tokens through the PKCS#11 URI.

(BZ#1668717)

5.4.9. Compilers and development tools

The **glibc** dynamic loader no longer fails while loading a shared library that uses **DT_FILTER** and has a constructor

Prior to this update, a defect in the dynamic loader implementation of shared objects as filters caused the dynamic loader to fail while loading a shared library that uses a filter and has a constructor. With this release, the dynamic loader implementation of filters (**DT_FILTER**) has been fixed to correctly handle such shared libraries. As a result, the dynamic loader now works as expected in the mentioned scenario.

(BZ#1812756)

glibc can now remove pseudo-mounts from the `getmntent()` list

The kernel includes **automount** pseudo-entries in the tables exposed to userspace. Consequently, programs that use the **getmntent()** API see both regular mounts and these pseudo-mounts in the list. The pseudo-mounts do not correspond to real mounts, nor include valid information.

With this update, if the mount entry has the **ignore** mount option present in the **automount(8)** configuration the **glibc** library now removes these pseudo-mounts from the **getmntent()** list. Programs that expect the previous behavior have to use a different API.

(BZ#1743445)

The `movv1qi` pattern no longer causes miscompilation in the auto-vectorized code on IBM Z

Prior to this update, wrong load instructions were emitted for the **movv1qi** pattern. As a consequence, when auto-vectorization was in effect, a miscompilation could occur on IBM Z systems. This update fixes the **movv1qi** pattern, and as a result, code compiles and runs correctly now.

(BZ#1784758)

PAPI_event_name_to_code() now works correctly in multiple threads

Prior to this update, the PAPI internal code did not handle thread coordination properly. As a consequence, when multiple threads used the **PAPI_event_name_to_code()** operation, a race condition occurred and the operation failed. This update enhances the handling of multiple threads in the PAPI internal code. As a result, multithreaded code using the **PAPI_event_name_to_code()** operation now works correctly.

(BZ#1807346)

Improved performance for the `glibc` math functions on IBM Power Systems

Previously, the **glibc** math functions performed unnecessary floating point status updates and system calls on IBM Power Systems, which negatively affected the performance. This update removes the unnecessary floating point status update, and improves the implementations of: **ceil()**, **ceilf()**, **fegetmode()**, **fesetmode()**, **fesetenv()**, **fegetexcept()**, **feenableexcept()**, **fedisableexcept()**, **fegetround()** and **fesetround()**. As a result, the performance of the math library is improved on IBM Power Systems.

(BZ#1783303)

Memory protection keys are now supported on IBM Power

On IBM Power Systems, the memory protection key interfaces **pkey_set** and **pkey_get** were previously stub functions, and consequently always failed. This update implements the interfaces, and as a result, the GNU C Library (**glibc**) now supports memory protection keys on IBM Power Systems.

Note that memory protection keys currently require the hash-based memory management unit (MMU), therefore you might have to boot certain systems with the **disable_radix** kernel parameter.

(BZ#1642150)

papi-testsuite and papi-devel now install the required papi-libs package

Previously, the **papi-testsuite** and **papi-devel** RPM packages did not declare a dependency on the matching **papi-libs** package. Consequently, the tests failed to run, and developers did not have the required version of the **papi** shared library available for their applications.

With this update, when the user installs either the **papi-testsuite** or **papi-devel** packages, the **papi-libs**

package is also installed. As a result, the **papi-testsuite** now has the correct library allowing the tests to run, and developers using **papi-devel** have their executables linked with the appropriate version of the **papi** shared library.

([BZ#1664056](#))

Installing the **lldb** packages for multiple architectures no longer leads to file conflicts

Previously, the **lldb** packages installed architecture-dependent files in architecture-independent locations. As a consequence, installing both 32-bit and 64-bit versions of the packages led to file conflicts. This update packages the files in correct architecture-dependent locations. As a result, the installation of **lldb** in the described scenario completes successfully.

([BZ#1841073](#))

getaddrinfo now correctly handles a memory allocation failure

Previously, after a memory allocation failure, the **getaddrinfo** function of the GNU C Library **glibc** did not release the internal resolver context. As a consequence, **getaddrinfo** was not able to reload the **/etc/resolv.conf** file for the rest of the lifetime of the calling thread, resulting in a possible memory leak.

This update modifies the error handling path with an additional release operation for the resolver context. As a result, **getaddrinfo** reloads **/etc/resolv.conf** with new configuration values even after an intermittent memory allocation failure.

([BZ#1810146](#))

glibc avoids certain failures caused by IFUNC resolver ordering

Previously, the implementation of the **librt** and **libpthread** libraries of the GNU C Library **glibc** contained the indirect function (IFUNC) resolvers for the following functions: **clock_gettime**, **clock_getcpuclockid**, **clock_nanosleep**, **clock_settime**, **vfork**. In some cases, the IFUNC resolvers could execute before the **librt** and **libpthread** libraries were relocated. Consequently, applications would fail in the **glibc** dynamic loader during early program startup.

With this release, the implementations of these functions have been moved into the **libc** component of **glibc**, which prevents the described problem from occurring.

([BZ#1748197](#))

Assertion failures no longer occur during **pthread_create**

Previously, the **glibc** dynamic loader did not roll back changes to the internal Thread Local Storage (TLS) module ID counter. As a consequence, an assertion failure in the **pthread_create** function could occur after the **dlopen** function had failed in certain ways. With this fix, the **glibc** dynamic loader updates the TLS module ID counter at a later point in time, after certain failures can no longer happen. As a result, the assertion failures no longer occur.

([BZ#1774115](#))

glibc now installs correct dependencies for 32-bit applications using **nss_db**

Previously, the **nss_db.x86_64** package did not declare dependencies on the **nss_db.i686** package. Therefore automated installation did not install **nss_db.i686** on the system, despite having a 32-bit environment **glibc.i686** installed. As a consequence, 32-bit applications using **nss_db** failed to perform accurate user database lookups, while 64-bit applications in the same setup worked correctly.

With this update, the **glibc** packages now have weak dependencies that trigger the installation of the

nss_db.i686 package when both **glibc.i686** and **nss_db** are installed on the system. As a result, 32-bit applications using **nss_db** now work correctly, even if the system administrator has not explicitly installed the **nss_db.i686** package.

([BZ#1807824](#))

glibc locale information updated with Odia language

The name of Indian state previously known as Orissa has changed to Odisha, and the name of its official language has changed from Oriya to Odia. With this update, the **glibc** locale information reflects the new name of the language.

([BZ#1757354](#))

LLVM sub packages now install arch-dependent files in arch-dependent locations

Previously, LLVM sub packages installed arch-dependent files in arch-independent locations. This resulted in conflicts when installing 32 and 64 bit versions of LLVM. With this update, package files are now correctly installed in arch-dependent locations, avoiding version conflicts.

([BZ#1820319](#))

Password and group lookups no longer fail in glibc

Previously, the **nss_compat** module of the **glibc** library overwrote the **errno** status with incorrect error codes during processing of password and group entries. Consequently, applications did not resize buffers as expected, causing password and group lookups to fail. This update fixes the problem, and the lookups now complete as expected.

([BZ#1836867](#))

5.4.10. Identity Management

SSSD no longer downloads every rule with a wildcard character by default

Previously, the **ldap_sudo_include_regexp** option was incorrectly set to **true** by default. As a consequence, when SSSD started running or after updating SSSD rules, SSSD downloaded every rule that contained a wildcard character (*) in the **sudoHost** attribute. This update fixes the bug, and the **ldap_sudo_include_regexp** option is now properly set to **false** by default. As a result, the described problem no longer occurs.

([BZ#1827615](#))

krb5 now only requests permitted encryption types

Previously, permitted encryption types specified in the **permitted_enctypes** variable in the **/etc/krb5.conf** file did not apply to the default encryption types if the **default_tgs_enctypes** or **default_tkt_enctypes** attributes were not set. Consequently, Kerberos clients were able to request deprecated cipher suites like RC4, which may cause other processes to fail. With this update, encryption types specified in the **permitted_enctypes** variable apply to the default encryption types as well, and only permitted encryption types are requested.

The RC4 cipher suite, which has been deprecated in RHEL 8, is the default encryption type for users, services, and trusts between Active Directory (AD) domains in an AD forest.

- To ensure support for strong AES encryption types between AD domains in an AD forest, see the [AD DS: Security: Kerberos "Unsupported etype" error when accessing a resource in a trusted domain](#) Microsoft article.

- To enable support for the deprecated RC4 encryption type in an IdM server for backwards compatibility with AD, use the **update-crypto-policies --set DEFAULT:AD-SUPPORT** command.

(BZ#1791062)

KDCs now correctly enforce password lifetime policy from LDAP backends

Previously, non-IPA Kerberos Distribution Centers (KDCs) did not ensure maximum password lifetimes because the Kerberos LDAP backend incorrectly enforced password policies. With this update, the Kerberos LDAP backend has been fixed, and password lifetimes behave as expected.

(BZ#1784655)

Password expiration notifications sent to AD clients using SSSD

Previously, Active Directory clients (non-IdM) using SSSD were not sent password expiration notices because of a recent change in the SSSD interface for acquiring Kerberos credentials.

The Kerberos interface has been updated and expiration notices are now sent correctly.

(BZ#1820311)

Directory Server no longer leaks memory when using indirect COS definitions

Previously, after processing an indirect Class Of Service (COS) definition, Directory Server leaked memory for each search operation that used an indirect COS definition. With this update, Directory Server frees all internal COS structures associated with the database entry after it has been processed. As a result, the server no longer leaks memory when using indirect COS definitions.

(BZ#1816862)

Adding ID overrides of AD users now works in IdM Web UI

Previously, adding ID overrides of Active Directory (AD) users to Identity Management (IdM) groups in the Default Trust View for the purpose of granting access to management roles failed when using the IdM Web UI. This update fixes the bug. As a result, you can now use both the Web UI as well as the IdM command-line interface (CLI) in this scenario.

(BZ#1651577)

FreeRADIUS no longer generates certificates during package installation

Previously, FreeRADIUS generated certificates during package installation, resulting in the following issues:

- If FreeRADIUS was installed using Kickstart, certificates might be generated at a time when entropy on the system was insufficient, resulting in either a failed installation or a less secure certificate.
- The package was difficult to build as part of an image, such as a container, because the package installation occurs on the builder machine instead of the target machine. All instances that are spawned from the image had the same certificate information.
- It was difficult for an end-user to generate a simple VM in their environment as the certificates would have to be removed and regenerated manually.

With this update, the FreeRADIUS installation no longer generates default self-signed CA certificates nor subordinate CA certificates. When FreeRADIUS is launched via **systemd**:

- If all of the required certificates are missing, a set of default certificates are generated.
- If one or more of the expected certificates are present, it does not generate new certificates.

([BZ#1672285](#))

FreeRADIUS now generates FIPS-compliant Diffie-Hellman parameters

Due to new FIPS requirements that do not allow **openssl** to generate Diffie-Hellman (dh) parameters via **dhparam**, the dh parameter generation has been removed from the FreeRADIUS bootstrap scripts and the file, **rfc3526-group-18-8192.dhparam**, is included with the FreeRADIUS packages for all systems, and thus enables FreeRADIUS to start in FIPS mode.

Note that you can customize **/etc/raddb/certs/bootstrap** and **/etc/raddb/certs/Makefile** to restore the DH parameter generation if required.

([BZ#1859527](#))

Updating Healthcheck now properly updates both **ipa-healthcheck-core** and **ipa-healthcheck**

Previously, entering **yum update healthcheck** did not update the **ipa-healthcheck** package but replaced it with the **ipa-healthcheck-core** package. As a consequence, the **ipa-healthcheck** command did not work after the update.

This update fixes the bug, and updating **ipa-healthcheck** now correctly updates both the **ipa-healthcheck** package and the **ipa-healthcheck-core** package. As a result, the **Healthcheck** tool works correctly after the update.

([BZ#1852244](#))

5.4.11. Graphics infrastructures

Laptops with hybrid Nvidia GPUs can now successfully resume from suspend

Previously, the **nouveau** graphics driver sometimes could not power on hybrid Nvidia GPUs on certain laptops from power-save mode. As a result, the laptops failed to resume from suspend.

With this update, several problems in the Runtime Power Management (**runpm**) system have been fixed. As a result, the laptops with hybrid graphics can now successfully resume from suspend.

([JIRA:RHELPLAN-57572](#))

5.4.12. Virtualization

Migrating virtual machines with the default CPU model now works more reliably

Previously, if a virtual machine (VM) was created without a specific CPU model, QEMU used a default model that was not visible to the **libvirt** service. As a consequence, it was possible to migrate the VM to a host that did not support the default CPU model of the VM, which sometimes caused crashes and incorrect behavior in the guest OS after the migration.

With this update, **libvirt** explicitly uses the **qemu64** model as default in the XML configuration of the VM. As a result, if the user attempts migrating a VM with the default CPU model to a host that does not support that model, **libvirt** correctly generates an error message.

Note, however, that Red Hat strongly recommends using a specific CPU model for your VMs.

(JIRA:RHELPLAN-45906)

5.4.13. Containers

Notes on FIPS support with Podman

The Federal Information Processing Standard (FIPS) requires certified modules to be used. Previously, Podman correctly installed certified modules in containers by enabling the proper flags at startup. However, in this release, Podman does not properly set up the additional application helpers normally provided by the system in the form of the FIPS system-wide crypto-policy. Although setting the system-wide crypto-policy is not required by the certified modules it does improve the ability of applications to use crypto modules in compliant ways. To work around this problem, change your container to run the **update-crypto-policies --set FIPS** command before any other application code was executed. The **update-crypto-policies --set FIPS** command is no longer required with this fix.

(BZ#1804193)

5.5. TECHNOLOGY PREVIEWS

This part provides a list of all Technology Previews available in Red Hat Enterprise Linux 8.3.

For information on Red Hat scope of support for Technology Preview features, see [Technology Preview Features Support Scope](#).

5.5.1. Networking

Enabled the **xt_u32** Netfilter module

The **xt_u32** Netfilter module is now available in the **kernel-modules-extra** rpm. This module helps in packet forwarding based on the data that is inaccessible to other protocol-based packet filters and thus eases manual migration to **nftables**. However, **xt_u32** Netfilter module is not supported by Red Hat.

(BZ#1834769)

nmstate available as a Technology Preview

Nmstate is a network API for hosts. The **nmstate** packages, available as a Technology Preview, provide a library and the **nmstatectl** command-line utility to manage host network settings in a declarative manner. The networking state is described by a pre-defined schema. Reporting of the current state and changes to the desired state both conform to the schema.

For further details, see the **/usr/share/doc/nmstate/README.md** file and the examples in the **/usr/share/doc/nmstate/examples** directory.

(BZ#1674456)

AF_XDP available as a Technology Preview

Address Family eXpress Data Path (AF_XDP) socket is designed for high-performance packet processing. It accompanies **XDP** and grants efficient redirection of programmatically selected packets to user space applications for further processing.

(BZ#1633143)

XDP available as a Technology Preview

The eXpress Data Path (XDP) feature, which is available as a Technology Preview, provides a means to

attach extended Berkeley Packet Filter (eBPF) programs for high-performance packet processing at an early point in the kernel ingress data path, allowing efficient programmable packet analysis, filtering, and manipulation.

(BZ#1503672)

KTLS available as a Technology Preview

In Red Hat Enterprise Linux 8, Kernel Transport Layer Security (KTLS) is provided as a Technology Preview. KTLS handles TLS records using the symmetric encryption or decryption algorithms in the kernel for the AES-GCM cipher. KTLS also provides the interface for offloading TLS record encryption to Network Interface Controllers (NICs) that support this functionality.

(BZ#1570255)

XDP features that are available as Technology Preview

Red Hat provides the usage of the following eXpress Data Path (XDP) features as unsupported Technology Preview:

- Loading XDP programs on architectures other than AMD and Intel 64-bit. Note that the **libxdp** library is not available for architectures other than AMD and Intel 64-bit.
- The **XDP_TX** and **XDP_REDIRECT** return codes.
- The XDP hardware offloading. Before using this feature, see [Unloading XDP programs on Netronome network cards that use the nfp driver fails](#).

(BZ#1889737)

act_mpls module available as a Technology Preview

The **act_mpls** module is now available in the **kernel-modules-extra** rpm as a Technology Preview. The module allows the application of Multiprotocol Label Switching (MPLS) actions with Traffic Control (TC) filters, for example, push and pop MPLS label stack entries with TC filters. The module also allows the Label, Traffic Class, Bottom of Stack, and Time to Live fields to be set independently.

(BZ#1839311)

Multipath TCP is now available as a Technology Preview

Multipath TCP (MPTCP), an extension to TCP, is now available as a Technology Preview. MPTCP improves resource usage within the network and resilience to network failure. For example, with Multipath TCP on the RHEL server, smartphones with MPTCP v1 enabled can connect to an application running on the server and switch between Wi-Fi and cellular networks without interrupting the connection to the server.

Note that either the applications running on the server must natively support MPTCP or administrators must load an **eBPF** program into the kernel to dynamically change **IPPROTO_TCP** to **IPPROTO_MPTCP**.

For further details see, [Getting started with Multipath TCP](#).

(JIRA:RHELPLAN-41549)

The systemd-resolved service is now available as a Technology Preview

The **systemd-resolved** service provides name resolution to local applications. The service implements a caching and validating DNS stub resolver, an Link-Local Multicast Name Resolution (LLMNR), and Multicast DNS resolver and responder.

Note that, even if the **systemd** package provides **systemd-resolved**, this service is an unsupported Technology Preview.

(BZ#1906489)

5.5.2. Kernel

The **kexec fast reboot** feature is available as Technology Preview

The **kexec fast reboot** feature continues to be available as a Technology Preview. **kexec fast reboot** significantly speeds the boot process by allowing the kernel to boot directly into the second kernel without passing through the Basic Input/Output System (BIOS) first. To use this feature:

1. Load the **kexec** kernel manually.
2. Reboot the operating system.

(BZ#1769727)

eBPF available as a Technology Preview

Extended Berkeley Packet Filter (eBPF) is an in-kernel virtual machine that allows code execution in the kernel space, in the restricted sandbox environment with access to a limited set of functions.

The virtual machine includes a new system call **bpf()**, which supports creating various types of maps, and also allows to load programs in a special assembly-like code. The code is then loaded to the kernel and translated to the native machine code with just-in-time compilation. Note that the **bpf()** syscall can be successfully used only by a user with the **CAP_SYS_ADMIN** capability, such as the root user. See the **bpf(2)** man page for more information.

The loaded programs can be attached onto a variety of points (sockets, tracepoints, packet reception) to receive and process data.

There are numerous components shipped by Red Hat that utilize the **eBPF** virtual machine. Each component is in a different development phase, and thus not all components are currently fully supported. All components are available as a Technology Preview, unless a specific component is indicated as supported.

The following notable **eBPF** components are currently available as a Technology Preview:

- **bpftrace**, a high-level tracing language that utilizes the **eBPF** virtual machine.
- **AF_XDP**, a socket for connecting the **eXpress Data Path (XDP)** path to user space for applications that prioritize packet processing performance.

(BZ#1559616)

The **igc** driver available as a Technology Preview for RHEL 8

The **igc** Intel 2.5G Ethernet Linux wired LAN driver is now available on all architectures for RHEL 8 as a Technology Preview. The **ethtool** utility also supports **igc** wired LANs.

(BZ#1495358)

Soft-RoCE available as a Technology Preview

Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) is a network protocol which implements RDMA over Ethernet. Soft-RoCE is the software implementation of RoCE which supports two protocol versions, RoCE v1 and RoCE v2. The Soft-RoCE driver, **rdma_rxe**, is available as an unsupported Technology Preview in RHEL 8.

(BZ#1605216)

5.5.3. File systems and storage

NVMe/TCP is available as a Technology Preview

Accessing and sharing Nonvolatile Memory Express (NVMe) storage over TCP/IP networks (NVMe/TCP) and its corresponding **nvme-tcp.ko** and **nvmet-tcp.ko** kernel modules have been added as a Technology Preview.

The use of NVMe/TCP as either a storage client or a target is manageable with tools provided by the **nvme-cli** and **nvmetcli** packages.

The NVMe/TCP target Technology Preview is included only for testing purposes and is not currently planned for full support.

(BZ#1696451)

File system DAX is now available for ext4 and XFS as a Technology Preview

In Red Hat Enterprise Linux 8, file system DAX is available as a Technology Preview. DAX provides a means for an application to directly map persistent memory into its address space. To use DAX, a system must have some form of persistent memory available, usually in the form of one or more Non-Volatile Dual In-line Memory Modules (NVDIMMs), and a file system that supports DAX must be created on the NVDIMM(s). Also, the file system must be mounted with the **dax** mount option. Then, an **mmap** of a file on the dax-mounted file system results in a direct mapping of storage into the application's address space.

(BZ#1627455)

OverlayFS

OverlayFS is a type of union file system. It enables you to overlay one file system on top of another. Changes are recorded in the upper file system, while the lower file system remains unmodified. This allows multiple users to share a file-system image, such as a container or a DVD-ROM, where the base image is on read-only media.

OverlayFS remains a Technology Preview under most circumstances. As such, the kernel logs warnings when this technology is activated.

Full support is available for OverlayFS when used with supported container engines (**podman**, **cri-o**, or **buildah**) under the following restrictions:

- OverlayFS is supported for use only as a container engine graph driver. Its use is supported only for container COW content, not for persistent storage. You must place any persistent storage on non-OverlayFS volumes. You can use only the default container engine configuration: one level of overlay, one lowerdir, and both lower and upper levels are on the same file system.
- Only XFS is currently supported for use as a lower layer file system.

Additionally, the following rules and limitations apply to using OverlayFS:

- The OverlayFS kernel ABI and user-space behavior are not considered stable, and might change in future updates.
- OverlayFS provides a restricted set of the POSIX standards. Test your application thoroughly before deploying it with OverlayFS. The following cases are not POSIX-compliant:
 - Lower files opened with **O_RDONLY** do not receive **st_atime** updates when the files are read.
 - Lower files opened with **O_RDONLY**, then mapped with **MAP_SHARED** are inconsistent with subsequent modification.
 - Fully compliant **st_ino** or **d_ino** values are not enabled by default on RHEL 8, but you can enable full POSIX compliance for them with a module option or mount option. To get consistent inode numbering, use the **xino=on** mount option.

You can also use the **redirect_dir=on** and **index=on** options to improve POSIX compliance. These two options make the format of the upper layer incompatible with an overlay without these options. That is, you might get unexpected results or errors if you create an overlay with **redirect_dir=on** or **index=on**, unmount the overlay, then mount the overlay without these options.

- To determine whether an existing XFS file system is eligible for use as an overlay, use the following command and see if the **ftype=1** option is enabled:

```
# xfs_info /mount-point | grep ftype
```

- SELinux security labels are enabled by default in all supported container engines with OverlayFS.
- Several known issues are associated with OverlayFS in this release. For details, see *Non-standard behavior* in the Linux kernel documentation:
<https://www.kernel.org/doc/Documentation/filesystems/overlayfs.txt>.

For more information about OverlayFS, see the Linux kernel documentation:
<https://www.kernel.org/doc/Documentation/filesystems/overlayfs.txt>.

(BZ#1690207)

Stratis is now available as a Technology Preview

Stratis is a new local storage manager. It provides managed file systems on top of pools of storage with additional features to the user.

Stratis enables you to more easily perform storage tasks such as:

- Manage snapshots and thin provisioning
- Automatically grow file system sizes as needed
- Maintain file systems

To administer Stratis storage, use the **stratis** utility, which communicates with the **stratisd** background service.

Stratis is provided as a Technology Preview.

For more information, see the Stratis documentation: [Setting up Stratis file systems](#).

RHEL 8.3 updates Stratis to version 2.1.0. For more information, see [Stratis 2.1.0 Release Notes](#).

(JIRA:RHELPLAN-1212)

IdM now supports setting up a Samba server on an IdM domain member as a Technology Preview

With this update, you can now set up a Samba server on an Identity Management (IdM) domain member. The new **ipa-client-samba** utility provided by the same-named package adds a Samba-specific Kerberos service principal to IdM and prepares the IdM client. For example, the utility creates the **/etc/samba/smb.conf** with the ID mapping configuration for the **sss** ID mapping back end. As a result, administrators can now set up Samba on an IdM domain member.

Due to IdM Trust Controllers not supporting the Global Catalog Service, AD-enrolled Windows hosts cannot find IdM users and groups in Windows. Additionally, IdM Trust Controllers do not support resolving IdM groups using the Distributed Computing Environment / Remote Procedure Calls (DCE/RPC) protocols. As a consequence, AD users can only access the Samba shares and printers from IdM clients.

For details, see [Setting up Samba on an IdM domain member](#).

(JIRA:RHELPLAN-13195)

5.5.4. High availability and clusters

Local mode version of **pcs cluster setup** command available as a technology preview

By default, the **pcs cluster setup** command automatically synchronizes all configuration files to the cluster nodes. In Red Hat Enterprise Linux 8.3, the **pcs cluster setup** command provides the **--corosync-conf** option as a technology preview. Specifying this option switches the command to **local** mode. In this mode, **pcs** creates a **corosync.conf** file and saves it to a specified file on the local node only, without communicating with any other node. This allows you to create a **corosync.conf** file in a script and handle that file by means of the script.

([BZ#1839637](#))

Pacemaker **podman** bundles available as a Technology Preview

Pacemaker container bundles now run on the **podman** container platform, with the container bundle feature being available as a Technology Preview. There is one exception to this feature being Technology Preview: Red Hat fully supports the use of Pacemaker bundles for Red Hat Openstack.

([BZ#1619620](#))

Heuristics in **corosync-qdevice** available as a Technology Preview

Heuristics are a set of commands executed locally on startup, cluster membership change, successful connect to **corosync-qnetd**, and, optionally, on a periodic basis. When all commands finish successfully on time (their return error code is zero), heuristics have passed; otherwise, they have failed. The heuristics result is sent to **corosync-qnetd** where it is used in calculations to determine which partition should be quorate.

([BZ#1784200](#))

New **fence-agents-heuristics-ping** fence agent

As a Technology Preview, Pacemaker now supports the **fence_heuristics_ping** agent. This agent aims to open a class of experimental fence agents that do no actual fencing by themselves but instead exploit the behavior of fencing levels in a new way.

If the heuristics agent is configured on the same fencing level as the fence agent that does the actual fencing but is configured before that agent in sequence, fencing issues an **off** action on the heuristics agent before it attempts to do so on the agent that does the fencing. If the heuristics agent gives a negative result for the **off** action it is already clear that the fencing level is not going to succeed, causing Pacemaker fencing to skip the step of issuing the **off** action on the agent that does the fencing. A heuristics agent can exploit this behavior to prevent the agent that does the actual fencing from fencing a node under certain conditions.

A user might want to use this agent, especially in a two-node cluster, when it would not make sense for a node to fence the peer if it can know beforehand that it would not be able to take over the services properly. For example, it might not make sense for a node to take over services if it has problems reaching the networking uplink, making the services unreachable to clients, a situation which a ping to a router might detect in that case.

(BZ#1775847)

5.5.5. Identity Management

Identity Management JSON-RPC API available as Technology Preview

An API is available for Identity Management (IdM). To view the API, IdM also provides an API browser as Technology Preview.

In Red Hat Enterprise Linux 7.3, the IdM API was enhanced to enable multiple versions of API commands. Previously, enhancements could change the behavior of a command in an incompatible way. Users are now able to continue using existing tools and scripts even if the IdM API changes. This enables:

- Administrators to use previous or later versions of IdM on the server than on the managing client.
- Developers to use a specific version of an IdM call, even if the IdM version changes on the server.

In all cases, the communication with the server is possible, regardless if one side uses, for example, a newer version that introduces new options for a feature.

For details on using the API, see [Using the Identity Management API to Communicate with the IdM Server \(TECHNOLOGY PREVIEW\)](#).

(BZ#1664719)

DNSSEC available as Technology Preview in IdM

Identity Management (IdM) servers with integrated DNS now support DNS Security Extensions (DNSSEC), a set of extensions to DNS that enhance security of the DNS protocol. DNS zones hosted on IdM servers can be automatically signed using DNSSEC. The cryptographic keys are automatically generated and rotated.

Users who decide to secure their DNS zones with DNSSEC are advised to read and follow these documents:

- DNSSEC Operational Practices, Version 2: <http://tools.ietf.org/html/rfc6781#section-2>

- Secure Domain Name System (DNS) Deployment Guide: <http://dx.doi.org/10.6028/NIST.SP.800-81-2>
- DNSSEC Key Rollover Timing Considerations: <http://tools.ietf.org/html/rfc7583>

Note that IdM servers with integrated DNS use DNSSEC to validate DNS answers obtained from other DNS servers. This might affect the availability of DNS zones that are not configured in accordance with recommended naming practices.

([BZ#1664718](#))

5.5.6. Desktop

GNOME for the 64-bit ARM architecture available as a Technology Preview

The GNOME desktop environment is now available for the 64-bit ARM architecture as a Technology Preview. This enables administrators to configure and manage servers from a graphical user interface (GUI) remotely, using the VNC session.

As a consequence, new administration applications are available on the 64-bit ARM architecture. For example: **Disk Usage Analyzer (baobab)**, **Firewall Configuration (firewall-config)**, **Red Hat Subscription Manager (subscription-manager)**, or the **Firefox** web browser. Using **Firefox**, administrators can connect to the local Cockpit daemon remotely.

(JIRA:RHELPLAN-27394, [BZ#1667225](#), [BZ#1667516](#), [BZ#1724302](#))

GNOME desktop on IBM Z is available as a Technology Preview

The GNOME desktop, including the Firefox web browser, is now available as a Technology Preview on the IBM Z architecture. You can now connect to a remote graphical session running GNOME using VNC to configure and manage your IBM Z servers.

(JIRA:RHELPLAN-27737)

5.5.7. Graphics infrastructures

VNC remote console available as a Technology Preview for the 64-bit ARM architecture

On the 64-bit ARM architecture, the Virtual Network Computing (VNC) remote console is available as a Technology Preview. Note that the rest of the graphics stack is currently unverified for the 64-bit ARM architecture.

([BZ#1698565](#))

Intel Tiger Lake graphics available as a Technology Preview

Intel Tiger Lake UP3 and UP4 Xe graphics are now available as a Technology Preview.

To enable hardware acceleration with Intel Tiger Lake graphics, add the following option on the kernel command line:

```
i915.force_probe=pci-id
```

In this option, replace *pci-id* with one of the following:

- The PCI ID of your Intel GPU

- The `*` character to enable the **i915** driver with all alpha-quality hardware

(BZ#1783396)

5.5.8. Red Hat Enterprise Linux System Roles

The postfix role of RHEL System Roles available as a Technology Preview

Red Hat Enterprise Linux System Roles provides a configuration interface for Red Hat Enterprise Linux subsystems, which makes system configuration easier through the inclusion of Ansible Roles. This interface enables managing system configurations across multiple versions of Red Hat Enterprise Linux, as well as adopting new major releases.

The **rhel-system-roles** packages are distributed through the AppStream repository.

The **postfix** role is available as a Technology Preview.

The following roles are fully supported:

- **kdump**
- **network**
- **selinux**
- **storage**
- **timesync**

For more information, see the Knowledgebase article about [RHEL System Roles](#).

(BZ#1812552)

5.5.9. Virtualization

KVM virtualization is usable in RHEL 8 Hyper-V virtual machines

As a Technology Preview, nested KVM virtualization can now be used on the Microsoft Hyper-V hypervisor. As a result, you can create virtual machines on a RHEL 8 guest system running on a Hyper-V host.

Note that currently, this feature only works on Intel systems. In addition, nested virtualization is in some cases not enabled by default on Hyper-V. To enable it, see the following Microsoft documentation:

<https://docs.microsoft.com/en-us/virtualization/hyper-v-on-windows/user-guide/nested-virtualization>

(BZ#1519039)

AMD SEV for KVM virtual machines

As a Technology Preview, RHEL 8 introduces the Secure Encrypted Virtualization (SEV) feature for AMD EPYC host machines that use the KVM hypervisor. If enabled on a virtual machine (VM), SEV encrypts VM memory so that the host cannot access data on the VM. This increases the security of the VM if the host is successfully infected by malware.

Note that the number of VMs that can use this feature at a time on a single host is determined by the host hardware. Current AMD EPYC processors support up to 509 running VMs using SEV.

Also note that for VMs with SEV configured to be able to boot, you must also configure the VM with a hard memory limit. To do so, add the following to the VM's XML configuration:

```
<memtune>
<hard_limit unit='KiB'>N</hard_limit>
</memtune>
```

The recommended value for N is equal to or greater than the guest RAM + 256 MiB. For example, if the guest is assigned 2 GiB RAM, N should be 2359296 or greater.

(BZ#1501618, BZ#1501607, JIRA:RHELPLAN-7677)

Intel vGPU

As a Technology Preview, it is now possible to divide a physical Intel GPU device into multiple virtual devices referred to as **mediated devices**. These mediated devices can then be assigned to multiple virtual machines (VMs) as virtual GPUs. As a result, these VMs share the performance of a single physical Intel GPU.

Note that only selected Intel GPUs are compatible with the vGPU feature. In addition, assigning a physical GPU to VMs makes it impossible for the host to use the GPU, and may prevent graphical display output on the host from working.

(BZ#1528684)

Creating nested virtual machines

Nested KVM virtualization is provided as a Technology Preview for KVM virtual machines (VMs) running on AMD64 and IBM Z systems hosts with RHEL 8. With this feature, a RHEL 7 or RHEL 8 VM that runs on a physical RHEL 8 host can act as a hypervisor, and host its own VMs.

Note that in RHEL 8.2 and later, nested virtualization is fully supported for VMs running on an Intel 64 host.

(JIRA:RHELPLAN-14047, JIRA:RHELPLAN-24437)

Select Intel network adapters now support SR-IOV in RHEL guests on Hyper-V

As a Technology Preview, Red Hat Enterprise Linux guest operating systems running on a Hyper-V hypervisor can now use the single-root I/O virtualization (SR-IOV) feature for Intel network adapters supported by the **ixgbevf** and **iavf** drivers. This feature is enabled when the following conditions are met:

- SR-IOV support is enabled for the network interface controller (NIC)
- SR-IOV support is enabled for the virtual NIC
- SR-IOV support is enabled for the virtual switch
- The virtual function (VF) from the NIC is attached to the virtual machine

The feature is currently supported with Microsoft Windows Server 2019 and 2016.

(BZ#1348508)

5.5.10. Containers

podman container image is available as a Technology Preview

The **registry.redhat.io/rhel8/podman** container image is a containerized implementation of the **podman** package. The **podman** tool is used for managing containers and images, volumes mounted into those containers, and pods made from groups of containers. Podman is based on the **libpod** library for container lifecycle management. The **libpod** library provides APIs for managing containers, pods, container images, and volumes. This container image allows create, modify and run container images without the need to install the **podman** package on your system. The use-case does not cover running this image in rootless mode as a non-root user. To pull the **registry.redhat.io/rhel8/podman** container image, you need an active Red Hat Enterprise Linux subscription.

([BZ#1627899](#))

crun is available as a Technology Preview

The **crun** OCI runtime has been added to the **container-roots:rhel8** module. The **crun** provides an access to run with cgroupsV2. The **crun** supports an annotation that allows the container to access the rootless users additional groups. This is useful for volume mounting in a directory that the user only have group access to, or the directory is setgid on it.

([BZ#1841438](#))

5.6. DEPRECATED FUNCTIONALITY

This part provides an overview of functionality that has been *deprecated* in Red Hat Enterprise Linux 8.

Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments. For the most recent list of deprecated functionality within a particular major release, refer to the latest version of release documentation.

The support status of deprecated functionality remains unchanged within Red Hat Enterprise Linux 8. For information about the length of support, see [Red Hat Enterprise Linux Life Cycle](#) and [Red Hat Enterprise Linux Application Streams Life Cycle](#).

Deprecated hardware components are not recommended for new deployments on the current or future major releases. Hardware driver updates are limited to security and critical fixes only. Red Hat recommends replacing this hardware as soon as reasonably feasible.

A package can be deprecated and not recommended for further use. Under certain circumstances, a package can be removed from a product. Product documentation then identifies more recent packages that offer functionality similar, identical, or more advanced to the one deprecated, and provides further recommendations.

For information regarding functionality that is present in RHEL 7 but has been *removed* in RHEL 8, see [Considerations in adopting RHEL 8](#) .

5.6.1. Installer and image creation

Several Kickstart commands and options have been deprecated

Using the following commands and options in RHEL 8 Kickstart files will print a warning in the logs.

- **auth** or **authconfig**
- **device**
- **deviceprobe**

- **dmraid**
- **install**
- **lilo**
- **lilocheck**
- **mouse**
- **multipath**
- **bootloader --upgrade**
- **ignoredisk --interactive**
- **partition --active**
- **reboot --kexec**

Where only specific options are listed, the base command and its other options are still available and not deprecated.

For more details and related changes in Kickstart, see the [Kickstart changes](#) section of the *Considerations in adopting RHEL 8* document.

(BZ#1642765)

The **--interactive** option of the **ignoredisk** Kickstart command has been deprecated

Using the **--interactive option** in future releases of Red Hat Enterprise Linux will result in a fatal installation error. It is recommended that you modify your Kickstart file to remove the option.

(BZ#1637872)

lorax-composer back end for Image Builder is deprecated in RHEL 8

The previous back end **lorax-composer** for Image Builder is considered deprecated. It will only receive select fixes for the rest of the Red Hat Enterprise Linux 8 life cycle and will be omitted from future major releases. Red Hat recommends that you uninstall **lorax-composer** and install **osbuild-composer** back end instead.

See [Composing a customized RHEL system image](#) for more details.

(BZ#1893767)

5.6.2. Software management

rpmbuild --sign is deprecated

With this update, the **rpmbuild --sign** command has become deprecated. Using this command in future releases of Red Hat Enterprise Linux can result in an error. It is recommended that you use the **rpmsign** command instead.

(BZ#1688849)

5.6.3. Shells and command-line tools

Metalink support for curl has been disabled.

A flaw was found in curl functionality in the way it handles credentials and file hash mismatch for content downloaded using the Metalink. This flaw allows malicious actors controlling a hosting server to:

- Trick users into downloading malicious content
- Gain unauthorized access to provided credentials without the user's knowledge

The highest threat from this vulnerability is confidentiality and integrity. To avoid this, the Metalink support for curl has been disabled from Red Hat Enterprise Linux 8.2.0.z.

As a workaround, execute the following command, after the Metalink file is downloaded:

```
wget --trust-server-names --input-metalink`
```

For example:

```
wget --trust-server-names --input-metalink <(curl -s $URL)
```

(BZ#1999620)

5.6.4. Infrastructure services

mailman is deprecated

With this update, the **mailman** packages have been marked as deprecated and will not be available in the future major releases of Red Hat Enterprise Linux.

(BZ#1890976)

5.6.5. Security

NSS SEED ciphers are deprecated

The Mozilla Network Security Services (**NSS**) library will not support TLS cipher suites that use a SEED cipher in a future release. To ensure smooth transition of deployments that rely on SEED ciphers when NSS removes support, Red Hat recommends enabling support for other cipher suites.

Note that SEED ciphers are already disabled by default in RHEL.

([BZ#1817533](#))

TLS 1.0 and TLS 1.1 are deprecated

The TLS 1.0 and TLS 1.1 protocols are disabled in the **DEFAULT** system-wide cryptographic policy level. If your scenario, for example, a video conferencing application in the Firefox web browser, requires using the deprecated protocols, switch the system-wide cryptographic policy to the **LEGACY** level:

```
# update-crypto-policies --set LEGACY
```

For more information, see the [Strong crypto defaults in RHEL 8 and deprecation of weak crypto algorithms](#) Knowledgebase article on the Red Hat Customer Portal and the **update-crypto-policies(8)** man page.

([BZ#1660839](#))

DSA is deprecated in RHEL 8

The Digital Signature Algorithm (DSA) is considered deprecated in Red Hat Enterprise Linux 8. Authentication mechanisms that depend on DSA keys do not work in the default configuration. Note that **OpenSSH** clients do not accept DSA host keys even in the **LEGACY** system-wide cryptographic policy level.

(BZ#1646541)

SSL2 Client Hello has been deprecated in NSS

The Transport Layer Security (**TLS**) protocol version 1.2 and earlier allow to start a negotiation with a **Client Hello** message formatted in a way that is backward compatible with the Secure Sockets Layer (**SSL**) protocol version 2. Support for this feature in the Network Security Services (**NSS**) library has been deprecated and it is disabled by default.

Applications that require support for this feature need to use the new **SSL_ENABLE_V2_COMPATIBLE_HELLO** API to enable it. Support for this feature may be removed completely in future releases of Red Hat Enterprise Linux 8.

(BZ#1645153)

TPM 1.2 is deprecated

The Trusted Platform Module (TPM) secure cryptoprocessor standard version was updated to version 2.0 in 2016. TPM 2.0 provides many improvements over TPM 1.2, and it is not backward compatible with the previous version. TPM 1.2 is deprecated in RHEL 8, and it might be removed in the next major release.

(BZ#1657927)

5.6.6. Networking

Network scripts are deprecated in RHEL 8

Network scripts are deprecated in Red Hat Enterprise Linux 8 and they are no longer provided by default. The basic installation provides a new version of the **ifup** and **ifdown** scripts which call the **NetworkManager** service through the **nmcli** tool. In Red Hat Enterprise Linux 8, to run the **ifup** and the **ifdown** scripts, NetworkManager must be running.

Note that custom commands in **/sbin/ifup-local**, **ifdown-pre-local** and **ifdown-local** scripts are not executed.

If any of these scripts are required, the installation of the deprecated network scripts in the system is still possible with the following command:

```
~]# yum install network-scripts
```

The **ifup** and **ifdown** scripts link to the installed legacy network scripts.

Calling the legacy network scripts shows a warning about their deprecation.

(BZ#1647725)

5.6.7. Kernel

Installing RHEL for Real Time 8 using diskless boot is now deprecated

Diskless booting allows multiple systems to share a root file system via the network. While convenient, diskless boot is prone to introducing network latency in realtime workloads. With a future minor update of RHEL for Real Time 8, the diskless booting feature will no longer be supported.

([BZ#1748980](#))

The **qla3xxx** driver is deprecated

The **qla3xxx** driver has been deprecated in RHEL 8. The driver will likely not be supported in future major releases of this product, and thus it is not recommended for new deployments.

([BZ#1658840](#))

The **dl2k**, **dnet**, **ethoc**, and **dlci** drivers are deprecated

The **dl2k**, **dnet**, **ethoc**, and **dlci** drivers have been deprecated in RHEL 8. The drivers will likely not be supported in future major releases of this product, and thus they are not recommended for new deployments.

([BZ#1660627](#))

The **rdma_rxe** Soft-RoCE driver is deprecated

Software Remote Direct Memory Access over Converged Ethernet (Soft-RoCE), also known as RXE, is a feature that emulates Remote Direct Memory Access (RDMA). In RHEL 8, the Soft-RoCE feature is available as an unsupported Technology Preview. However, due to stability issues, this feature has been deprecated and will be removed in RHEL 9.

([BZ#1878207](#))

5.6.8. File systems and storage

The **elevator** kernel command line parameter is deprecated

The **elevator** kernel command line parameter was used in earlier RHEL releases to set the disk scheduler for all devices. In RHEL 8, the parameter is deprecated.

The upstream Linux kernel has removed support for the **elevator** parameter, but it is still available in RHEL 8 for compatibility reasons.

Note that the kernel selects a default disk scheduler based on the type of device. This is typically the optimal setting. If you require a different scheduler, Red Hat recommends that you use **udev** rules or the Tuned service to configure it. Match the selected devices and switch the scheduler only for those devices.

For more information, see [Setting the disk scheduler](#).

([BZ#1665295](#))

LVM **mirror** is deprecated

The LVM **mirror** segment type is now deprecated. Support for **mirror** will be removed in a future major release of RHEL.

Red Hat recommends that you use LVM RAID 1 devices with a segment type of **raid1** instead of **mirror**. The **raid1** segment type is the default RAID configuration type and replaces **mirror** as the recommended solution.

To convert **mirror** devices to **raid1**, see [Converting a mirrored LVM device to a RAID1 logical volume](#) .

LVM **mirror** has several known issues. For details, see [known issues in file systems and storage](#) .

(BZ#1827628)

peripety is deprecated

The **peripety** package is deprecated since RHEL 8.3.

The Peripety storage event notification daemon parses system storage logs into structured storage events. It helps you investigate storage issues.

(BZ#1871953)

NFSv3 over UDP has been disabled

The NFS server no longer opens or listens on a User Datagram Protocol (UDP) socket by default. This change affects only NFS version 3 because version 4 requires the Transmission Control Protocol (TCP).

NFS over UDP is no longer supported in RHEL 8.

(BZ#1592011)

cramfs has been deprecated

Due to lack of users, the **cramfs** kernel module is deprecated. **squashfs** is recommended as an alternative solution.

(BZ#1794513)

5.6.9. Identity Management

openssh-ldap has been deprecated

The **openssh-ldap** subpackage has been deprecated in Red Hat Enterprise Linux 8 and will be removed in RHEL 9. As the **openssh-ldap** subpackage is not maintained upstream, Red Hat recommends using SSSD and the **sss_ssh_authorizedkeys** helper, which integrate better with other IdM solutions and are more secure.

By default, the SSSD **ldap** and **ipa** providers read the **sshPublicKey** LDAP attribute of the user object, if available. Note that you cannot use the default SSSD configuration for the **ad** provider or IdM trusted domains to retrieve SSH public keys from Active Directory (AD), since AD does not have a default LDAP attribute to store a public key.

To allow the **sss_ssh_authorizedkeys** helper to get the key from SSSD, enable the **ssh** responder by adding **ssh** to the **services** option in the **sssd.conf** file. See the **sssd.conf(5)** man page for details.

To allow **sshd** to use **sss_ssh_authorizedkeys**, add the **AuthorizedKeysCommand** **/usr/bin/sss_ssh_authorizedkeys** and **AuthorizedKeysCommandUser nobody** options to the **/etc/ssh/sshd_config** file as described by the **sss_ssh_authorizedkeys(1)** man page.

(BZ#1871025)

DES and 3DES encryption types have been removed

Due to security reasons, the Data Encryption Standard (DES) algorithm has been deprecated and disabled by default since RHEL 7. With the recent rebase of Kerberos packages, single-DES (DES) and triple-DES (3DES) encryption types have been removed from RHEL 8.

If you have configured services or users to only use DES or 3DES encryption, you might experience service interruptions such as:

- Kerberos authentication errors
- **unknown enctype** encryption errors
- Kerberos Distribution Centers (KDCs) with DES-encrypted Database Master Keys (**K/M**) fail to start

Perform the following actions to prepare for the upgrade:

1. Check if your KDC uses DES or 3DES encryption with the **krb5check** open source Python scripts. See [krb5check](#) on GitHub.
2. If you are using DES or 3DES encryption with any Kerberos principals, re-key them with a supported encryption type, such as Advanced Encryption Standard (AES). For instructions on re-keying, see [Retiring DES](#) from MIT Kerberos Documentation.
3. Test independence from DES and 3DES by temporarily setting the following Kerberos options before upgrading:
 - a. In **/var/kerberos/krb5kdc/kdc.conf** on the KDC, set **supported_enctypes** and do not include **des** or **des3**.
 - b. For every host, in **/etc/krb5.conf** and any files in **/etc/krb5.conf.d**, set **allow_weak_crypto** to **false**. It is false by default.
 - c. For every host, in **/etc/krb5.conf** and any files in **/etc/krb5.conf.d**, set **permitted_enctypes**, **default_tgs_enctypes**, and **default_tkt_enctypes** and do not include **des** or **des3**.
4. If you do not experience any service interruptions with the test Kerberos settings from the previous step, remove them and upgrade. You do not need those settings after upgrading to the latest Kerberos packages.

([BZ#1877991](#))

5.6.10. Desktop

The libgnome-keyring library has been deprecated

The **libgnome-keyring** library has been deprecated in favor of the **libsecret** library, as **libgnome-keyring** is not maintained upstream, and does not follow the necessary cryptographic policies for RHEL. The new **libsecret** library is the replacement that follows the necessary security standards.

([BZ#1607766](#))

The AlternateTab extension has been removed

The **gnome-shell-extension-alternate-tab** package, which provides the **AlternateTab** GNOME Shell extension, has been removed.

To configure the window-switching behavior, set a keyboard shortcut in keyboard settings. For more information, see the following article: [Using Alternate-Tab in Gnome 3.32 or later](#) . ([BZ#1922488](#))

5.6.11. Graphics infrastructures

AGP graphics cards are no longer supported

Graphics cards using the Accelerated Graphics Port (AGP) bus are not supported in Red Hat Enterprise Linux 8. Use the graphics cards with PCI-Express bus as the recommended replacement.

(BZ#1569610)

5.6.12. The web console

The web console no longer supports incomplete translations

The RHEL web console no longer provides translations for languages that have translations available for less than 50 % of the Console's translatable strings. If the browser requests translation to such a language, the user interface will be in English instead.

(BZ#1666722)

5.6.13. Red Hat Enterprise Linux System Roles

The **geoipupdate** package has been deprecated

The **geoipupdate** package requires a third-party subscription and it also downloads proprietary content. Therefore, the **geoipupdate** package has been deprecated, and will be removed in the next major RHEL version.

(BZ#1874892)

5.6.14. Virtualization

virt-manager has been deprecated

The Virtual Machine Manager application, also known as **virt-manager**, has been deprecated. The RHEL 8 web console, also known as **Cockpit**, is intended to become its replacement in a subsequent release. It is, therefore, recommended that you use the web console for managing virtualization in a GUI. Note, however, that some features available in **virt-manager** may not be yet available the RHEL 8 web console.

(JIRA:RHELPLAN-10304)

Virtual machine snapshots are not properly supported in RHEL 8

The current mechanism of creating virtual machine (VM) snapshots has been deprecated, as it is not working reliably. As a consequence, it is recommended not to use VM snapshots in RHEL 8.

Note that a new VM snapshot mechanism is under development and will be fully implemented in a future minor release of RHEL 8.

(BZ#1686057)

The Cirrus VGA virtual GPU type has been deprecated

With a future major update of Red Hat Enterprise Linux, the **Cirrus VGA** GPU device will no longer be supported in KVM virtual machines. Therefore, Red Hat recommends using the **stdvga**, **virtio-vga**, or **qxl** devices instead of Cirrus VGA.

(BZ#1651994)

SPICE has been deprecated

In RHEL 8.3, the SPICE remote display protocol has been deprecated. Note that SPICE will remain supported in RHEL 8, but Red Hat recommends using alternate solutions for remote display streaming:

- For remote console access, use the VNC protocol.
- For advanced remote display functions, use third party tools such as RDP, HP RGS, or Mechdyne TGX.

(BZ#1849563)

5.6.15. Containers

Podman varlink-based REST API V1 has been deprecated

The Podman varlink-based REST API V1 has been deprecated upstream in favor of the new Podman REST API V2. This functionality will be removed in a later release of Red Hat Enterprise Linux 8.

(JIRA:RHELPLAN-60226)

5.6.16. Deprecated packages

The following packages have been deprecated and will probably not be included in a future major release of Red Hat Enterprise Linux:

- 389-ds-base-legacy-tools
- authd
- custodia
- hostname
- libidn
- lorax-composer
- mercurial
- net-tools
- network-scripts
- nss-pam-ldapd
- sendmail
- yp-tools
- ypbind
- ypserv

5.7. KNOWN ISSUES

This part describes known issues in Red Hat Enterprise Linux 8.3.

5.7.1. Installer and image creation

The **auth** and **authconfig** Kickstart commands require the AppStream repository

The **authselect-compat** package is required by the **auth** and **authconfig** Kickstart commands during installation. Without this package, the installation fails if **auth** or **authconfig** are used. However, by design, the **authselect-compat** package is only available in the AppStream repository.

To work around this problem, verify that the BaseOS and AppStream repositories are available to the installer or use the **authselect** Kickstart command during installation.

(BZ#1640697)

The **reboot --kexec** and **inst.kexec** commands do not provide a predictable system state

Performing a RHEL installation with the **reboot --kexec** Kickstart command or the **inst.kexec** kernel boot parameters do not provide the same predictable system state as a full reboot. As a consequence, switching to the installed system without rebooting can produce unpredictable results.

Note that the **kexec** feature is deprecated and will be removed in a future release of Red Hat Enterprise Linux.

(BZ#1697896)

Network access is not enabled by default in the installation program

Several installation features require network access, for example, registration of a system using the Content Delivery Network (CDN), NTP server support, and network installation sources. However, network access is not enabled by default, and as a result, these features cannot be used until network access is enabled.

To work around this problem, add **ip=dhcp** to boot options to enable network access when the installation starts. Optionally, passing a Kickstart file or a repository located on the network using boot options also resolves the problem. As a result, the network-based installation features can be used.

(BZ#1757877)

The new **osbuild-composer** back end does not replicate the blueprint state from **lorax-composer** on upgrades

Image Builder users that are upgrading from the **lorax-composer** back end to the new **osbuild-composer** back end, blueprints can disappear. As a result, once the upgrade is complete, the blueprints do not display automatically. To work around this problem, perform the following steps.

Prerequisites

- You have the **composer-cli** CLI utility installed.

Procedure

1. Run the command to load the previous **lorax-composer** based blueprints into the new **osbuild-composer** back end:

```
$ for blueprint in $(find /var/lib/lorax/composer/blueprints/git/workspace/master -name
*.toml); do composer-cli blueprints push "${blueprint}"; done
```

-

As a result, the same blueprints are now available in **osbuild-composer** back end.

Additional resources

- For more details about this Known Issue, see the [Image Builder blueprints are no longer present following an update to Red Hat Enterprise Linux 8.3](#) article.

(BZ#1897383)

Self-signed HTTPS server cannot be used in Kickstart installation

Currently, the installer fails to install from a self-signed https server when the installation source is specified in the kickstart file and the **--noverifyssl** option is used:

```
url --url=https://SERVER/PATH --noverifyssl
```

To work around this problem, append the **inst.noverifyssl** parameter to the kernel command line when starting the kickstart installation.

For example:

```
inst.ks=<URL> inst.noverifyssl
```

(BZ#1745064)

GUI installation might fail if an attempt to unregister using the CDN is made before the repository refresh is completed

Since RHEL 8.2, when registering your system and attaching subscriptions using the Content Delivery Network (CDN), a refresh of the repository metadata is started by the GUI installation program. The refresh process is not part of the registration and subscription process, and as a consequence, the **Unregister** button is enabled in the **Connect to Red Hat** window. Depending on the network connection, the refresh process might take more than a minute to complete. If you click the **Unregister** button before the refresh process is completed, the GUI installation might fail as the unregister process removes the CDN repository files and the certificates required by the installation program to communicate with the CDN.

To work around this problem, complete the following steps in the GUI installation after you have clicked the **Register** button in the **Connect to Red Hat** window:

1. From the **Connect to Red Hat** window, click **Done** to return to the **Installation Summary** window.
2. From the **Installation Summary** window, verify that the **Installation Source** and **Software Selection** status messages in italics are not displaying any processing information.
3. When the Installation Source and Software Selection categories are ready, click **Connect to Red Hat**.
4. Click the **Unregister** button.

After performing these steps, you can safely unregister the system during the GUI installation.

(BZ#1821192)

Registration fails for user accounts that belong to multiple organizations

Currently, when you attempt to register a system with a user account that belongs to multiple organizations, the registration process fails with the error message **You must specify an organization for new units**.

To work around this problem, you can either:

- Use a different user account that does not belong to multiple organizations.
- Use the **Activation Key** authentication method available in the Connect to Red Hat feature for GUI and Kickstart installations.
- Skip the registration step in Connect to Red Hat and use Subscription Manager to register your system post-installation.

([BZ#1822880](#))

RHEL installer fails to start when InfiniBand network interfaces are configured using installer boot options

When you configure InfiniBand network interfaces at an early stage of RHEL installation using installer boot options (for example, to download installer image using PXE server), the installer fails to activate the network interfaces.

This issue occurs because the RHEL NetworkManager fails to recognize the network interfaces in InfiniBand mode, and instead configures Ethernet connections for the interfaces.

As a result, connection activation fails, and if the connectivity over the InfiniBand interface is required at an early stage, RHEL installer fails to start the installation.

To workaround this issue, create a new installation media including the updated Anaconda and NetworkManager packages, using the Lorax tool.

For more information about creating a new installation media including the updated Anaconda and NetworkManager packages, using the Lorax tool, see [Unable to install Red Hat Enterprise Linux 8.3.0 with InfiniBand network interfaces](#)

([BZ#1890261](#))

Anaconda installation fails when NVDIMM device namespace set to `devdax` mode.

Anaconda installation fails with a traceback after booting with NVDIMM device namespace set to **devdax** mode before the GUI installation.

To workaround this problem, reconfigure the NVDIMM device to set the namespace to a different mode than the **devdax** mode before the installation begins. As a result, you can proceed with the installation.

([BZ#1891827](#))

Local Media installation source is not detected when booting the installation from a USB that is created using a third party tool

When booting the RHEL installation from a USB that is created using a third party tool, the installer fails to detect the **Local Media** installation source (only 'Red Hat CDN' is detected).

This issue occurs because the default boot option **int.stage2=** attempts to search for **iso9660** image format. However, a third party tool might create an ISO image with a different format.

As a workaround, use either of the following solution:

- When booting the installation, click the **Tab** key to edit the kernel command line, and change the boot option **inst.stage2=** to **inst.repo=**.
- To create a bootable USB device on Windows, use Fedora Media Writer.
- When using a third party tool like Rufus to create a bootable USB device, first regenerate the RHEL ISO image on a Linux system, and then use the third party tool to create a bootable USB device.

For more information on the steps involved in performing any of the specified workaround, see, [Installation media is not auto detected during the installation of RHEL 8.3](#)

(BZ#1877697)

Anaconda now shows a dialog for **lidl** or unformatted DASD disks in text mode

Previously, during an installation in text mode, Anaconda failed to show a dialog for Linux disk layout (**lidl**) or unformatted Direct-Access Storage Device (DASD) disks. As a result, users were unable to utilize those disks for the installation.

With this update, in text mode Anaconda recognizes **lidl** and unformatted DASD disks and shows a dialog where users can format them properly for the future utilization for the installation.

(BZ#1874394)

Red Hat Insights client fails to register the operating system when using the graphical installer

Currently, the installation fails with an error at the end, which points to the Insights client.

To work around this problem, uncheck the **Connect to Red Hat Insights** option during the **Connect to Red Hat** step before registering the systems in the installer.

As a result, you can complete the installation and register to Insights afterwards by using this command:

```
# insights-client --register
```

(BZ#1931069)

5.7.2. Subscription management

syspurpose add-ons have no effect on the **subscription-manager attach --auto** output.

In Red Hat Enterprise Linux 8, four attributes of the **syspurpose** command-line tool have been added: **role**, **usage**, **service_level_agreement** and **add-ons**. Currently, only **role**, **usage** and **service_level_agreement** affect the output of running the **subscription-manager attach --auto** command. Users who attempt to set values to the **add-ons** argument will not observe any effect on the subscriptions that are auto-attached.

(BZ#1687900)

5.7.3. Infrastructure services

libmaxminddb-devel-debuginfo.rpm is removed when running **dnf update**

When performing the **dnf update** command, the binary **mmdblookup** tool is moved from the **libmaxminddb-devel** subpackage to the main **libmaxminddb** package. Consequently, the

libmaxminddb-devel-debuginfo.rpm is removed, which might create a broken update path for this package. To work around this problem, remove the **libmaxminddb-devel-debuginfo** prior to the execution of the **dnf update** command.

Note: **libmaxminddb-debuginfo** is the new **debuginfo** package.

(BZ#1642001)

5.7.4. Security

Users can run **sudo** commands as locked users

In systems where **sudoers** permissions are defined with the **ALL** keyword, **sudo** users with permissions can run **sudo** commands as users whose accounts are locked. Consequently, locked and expired accounts can still be used to execute commands.

To work around this problem, enable the newly implemented **runas_check_shell** option together with proper settings of valid shells in **/etc/shells**. This prevents attackers from running commands under system accounts such as **bin**.

(BZ#1786990)

GnuTLS fails to resume current session with the NSS server

When resuming a TLS (Transport Layer Security) 1.3 session, the **GnuTLS** client waits 60 milliseconds plus an estimated round trip time for the server to send session resumption data. If the server does not send the resumption data within this time, the client creates a new session instead of resuming the current session. This incurs no serious adverse effects except for a minor performance impact on a regular session negotiation.

(BZ#1677754)

libselinux-python is available only through its module

The **libselinux-python** package contains only Python 2 bindings for developing SELinux applications and it is used for backward compatibility. For this reason, **libselinux-python** is no longer available in the default RHEL 8 repositories through the **dnf install libselinux-python** command.

To work around this problem, enable both the **libselinux-python** and **python27** modules, and install the **libselinux-python** package and its dependencies with the following commands:

```
# dnf module enable libselinux-python
# dnf install libselinux-python
```

Alternatively, install **libselinux-python** using its install profile with a single command:

```
# dnf module install libselinux-python:2.8/common
```

As a result, you can install **libselinux-python** using the respective module.

(BZ#1666328)

udica processes UBI 8 containers only when started with **--env container=podman**

The Red Hat Universal Base Image 8 (UBI 8) containers set the **container** environment variable to the **oci** value instead of the **podman** value. This prevents the **udica** tool from analyzing a container JavaScript Object Notation (JSON) file.

To work around this problem, start a UBI 8 container using a **podman** command with the **--env container=podman** parameter. As a result, **udica** can generate an SELinux policy for a UBI 8 container only when you use the described workaround.

([BZ#1763210](#))

Negative effects of the default logging setup on performance

The default logging environment setup might consume 4 GB of memory or even more and adjustments of rate-limit values are complex when **systemd-journald** is running with **rsyslog**.

See the [Negative effects of the RHEL default logging setup on performance and their mitigations](#) Knowledgebase article for more information.

(JIRA:RHELPLAN-10431)

File permissions of **/etc/passwd-** are not aligned with the CIS RHEL 8 Benchmark 1.0.0

Because of an issue with the CIS Benchmark, the remediation of the SCAP rule that ensures permissions on the **/etc/passwd-** backup file configures permissions to **0644**. However, the **CIS Red Hat Enterprise Linux 8 Benchmark 1.0.0** requires file permissions **0600** for that file. As a consequence, the file permissions of **/etc/passwd-** are not aligned with the benchmark after remediation.

([BZ#1858866](#))

SELINUX=disabled in **/etc/selinux/config** does not work properly

Disabling SELinux using the **SELINUX=disabled** option in the **/etc/selinux/config** results in a process in which the kernel boots with SELinux enabled and switches to disabled mode later in the boot process. This might cause memory leaks.

To work around this problem, disable SELinux by adding the **selinux=0** parameter to the kernel command line as described in the [Changing SELinux modes at boot time](#) section of the [Using SELinux](#) title if your scenario really requires to completely disable SELinux.

(JIRA:RHELPLAN-34199)

ssh-keyscan cannot retrieve RSA keys of servers in FIPS mode

The **SHA-1** algorithm is disabled for RSA signatures in FIPS mode, which prevents the **ssh-keyscan** utility from retrieving RSA keys of servers operating in that mode.

To work around this problem, use ECDSA keys instead, or retrieve the keys locally from the **/etc/ssh/ssh_host_rsa_key.pub** file on the server.

([BZ#1744108](#))

OpenSSL incorrectly handles PKCS #11 tokens that does not support raw RSA or RSA-PSS signatures

The **OpenSSL** library does not detect key-related capabilities of PKCS #11 tokens. Consequently, establishing a TLS connection fails when a signature is created with a token that does not support raw RSA or RSA-PSS signatures.

To work around the problem, add the following lines after the **.include** line at the end of the **crypto_policy** section in the **/etc/pki/tls/openssl.cnf** file:


```
SignatureAlgorithms =
RSA+SHA256:RSA+SHA512:RSA+SHA384:ECDSA+SHA256:ECDSA+SHA512:ECDSA+SHA384
MaxProtocol = TLSv1.2
```

As a result, a TLS connection can be established in the described scenario.

(BZ#1685470)

OpenSSL in FIPS mode accepts only specific D-H parameters

In FIPS mode, Transport Security Layer (TLS) clients that use OpenSSL return a **bad dh value** error and abort TLS connections to servers that use manually generated parameters. This is because OpenSSL, when configured to work in compliance with FIPS 140-2, works only with D-H parameters compliant to NIST SP 800-56A rev3 Appendix D (groups 14, 15, 16, 17, and 18 defined in RFC 3526 and with groups defined in RFC 7919). Also, servers that use OpenSSL ignore all other parameters and instead select known parameters of similar size. To work around this problem, use only the compliant groups.

(BZ#1810911)

Removing the rpm-plugin-selinux package leads to removing all selinux-policy packages from the system

Removing the **rpm-plugin-selinux** package disables SELinux on the machine. It also removes all **selinux-policy** packages from the system. Repeated installation of the **rpm-plugin-selinux** package then installs the **selinux-policy-minimum** SELinux policy, even if the **selinux-policy-targeted** policy was previously present on the system. However, the repeated installation does not update the SELinux configuration file to account for the change in policy. As a consequence, SELinux is disabled even upon reinstallation of the **rpm-plugin-selinux** package.

To work around this problem:

1. Enter the **umount /sys/fs/selinux/** command.
2. Manually install the missing **selinux-policy-targeted** package.
3. Edit the **/etc/selinux/config** file so that the policy is equal to **SELINUX=enforcing**.
4. Enter the command **load_policy -i**.

As a result, SELinux is enabled and running the same policy as before.

(BZ#1641631)

systemd service cannot execute commands from arbitrary paths

The **systemd** service cannot execute commands from **/home/user/bin** arbitrary paths because the SELinux policy package does not include any such rule. Consequently, the custom services that are executed on non-system paths fail and eventually logs the Access Vector Cache (AVC) denial audit messages when SELinux denied access. To work around this problem, do one of the following:

- Execute the command using a **shell** script with the **-c** option. For example,

```
bash -c command
```

- Execute the command from a common path using **/bin**, **/sbin**, **/usr/sbin**, **/usr/local/bin**, and **/usr/local/sbin** common directories.

([BZ#1860443](#))

rpm_verify_permissions fails in the CIS profile

The **rpm_verify_permissions** rule compares file permissions to package default permissions. However, the Center for Internet Security (CIS) profile, which is provided by the **scap-security-guide** packages, changes some file permissions to be more strict than default. As a consequence, verification of certain files using **rpm_verify_permissions** fails.

To work around this problem, manually verify that these files have the following permissions:

- **/etc/cron.d** (0700)
- **/etc/cron.hourly** (0700)
- **/etc/cron.monthly** (0700)
- **/etc/crontab** (0600)
- **/etc/cron.weekly** (0700)
- **/etc/cron.daily** (0700)

([BZ#1843913](#))

Kickstart uses org_fedora_osc const instead of com_redhat_osc const in RHEL 8

The Kickstart references the Open Security Content Automation Protocol (OSCAP) Anaconda add-on as **org_fedora_osc const** instead of **com_redhat_osc const** which might cause confusion. That is done to preserve backward compatibility with Red Hat Enterprise Linux 7.

([BZ#1665082](#))

Certain sets of interdependent rules in SSG can fail

Remediation of **SCAP Security Guide** (SSG) rules in a benchmark can fail due to undefined ordering of rules and their dependencies. If two or more rules need to be executed in a particular order, for example, when one rule installs a component and another rule configures the same component, they can run in the wrong order and remediation reports an error. To work around this problem, run the remediation twice, and the second run fixes the dependent rules.

([BZ#1750755](#))

OSCAP Anaconda Addon does not install all packages in text mode

The **OSCAP Anaconda Addon** plugin cannot modify the list of packages selected for installation by the system installer if the installation is running in text mode. Consequently, when a security policy profile is specified using Kickstart and the installation is running in text mode, any additional packages required by the security policy are not installed during installation.

To work around this problem, either run the installation in graphical mode or specify all packages that are required by the security policy profile in the security policy in the **%packages** section in your Kickstart file.

As a result, packages that are required by the security policy profile are not installed during RHEL installation without one of the described workarounds, and the installed system is not compliant with the given security policy profile.

([BZ#1674001](#))

OSCAP Anaconda Addon does not correctly handle customized profiles

The **OSCAP Anaconda Addon** plugin does not properly handle security profiles with customizations in separate files. Consequently, the customized profile is not available in the RHEL graphical installation even when you properly specify it in the corresponding Kickstart section.

To work around this problem, follow the instructions in the [Creating a single SCAP data stream from an original DS and a tailoring file](#) Knowledgebase article. As a result of this workaround, you can use a customized SCAP profile in the RHEL graphical installation.

(BZ#1691305)

OSPP-based profiles are incompatible with GUI package groups.

GNOME packages installed by the *Server with GUI* package group require the **nfs-utils** package that is not compliant with the Operating System Protection Profile (OSPP). As a consequence, selecting the *Server with GUI* package group during the installation of a system with OSPP or OSPP-based profiles, for example, Security Technical Implementation Guide (STIG), OpenSCAP displays a warning that the selected package group is not compatible with the security policy. If the OSPP-based profile is applied after the installation, the system is not bootable. To work around this problem, do not install the *Server with GUI* package group or any other groups that install GUI when using the OSPP profile and OSPP-based profiles. When you use the *Server* or *Minimal Install* package groups instead, the system installs without issues and works correctly.

(BZ#1787156)

Installation with the Server with GUI or Workstation software selections and CIS security profile is not possible

The CIS security profile is not compatible with the **Server with GUI** and **Workstation** software selections. As a consequence, a RHEL 8 installation with the **Server with GUI** software selection and CIS profile is not possible. An attempted installation using the CIS profile and either of these software selections will generate the error message:

```
package xorg-x11-server-common has been added to the list of excluded packages, but it can't be removed from the current software selection without breaking the installation.
```

To work around the problem, do not use the CIS security profile with the **Server with GUI** or **Workstation** software selections.

(BZ#1843932)

Remediating service-related rules during kickstart installations might fail

During a kickstart installation, the OpenSCAP utility sometimes incorrectly shows that a service **enable** or **disable** state remediation is not needed. Consequently, OpenSCAP might set the services on the installed system to a non-compliant state. As a workaround, you can scan and remediate the system after the kickstart installation. This will fix the service-related issues.

(BZ#1834716)

Certain rsyslog priority strings do not work correctly

Support for the **GnuTLS** priority string for **imtcp** that allows fine-grained control over encryption is not complete. Consequently, the following priority strings do not work properly in **rsyslog**:

```
NONE:+VERS-ALL:-VERS-TLS1.3:+MAC-ALL:+DHE-RSA:+AES-256-GCM:+SIGN-RSA-SHA384:+COMP-ALL:+GROUP-ALL
```

■
To work around this problem, use only correctly working priority strings:

```
NONE:+VERS-ALL:-VERS-TLS1.3:+MAC-ALL:+ECDHE-RSA:+AES-128-CBC:+SIGN-RSA-SHA1:+COMP-ALL:+GROUP-ALL
```

As a result, current configurations must be limited to the strings that work correctly.

([BZ#1679512](#))

crypto-policies incorrectly allow Camellia ciphers

The RHEL 8 system-wide cryptographic policies should disable Camellia ciphers in all policy levels, as stated in the product documentation. However, the Kerberos protocol enables the ciphers by default.

To work around the problem, apply the **NO-CAMELLIA** subpolicy:

```
# update-crypto-policies --set DEFAULT:NO-CAMELLIA
```

In the previous command, replace **DEFAULT** with the cryptographic level name if you have switched from **DEFAULT** previously.

As a result, Camellia ciphers are correctly disallowed across all applications that use system-wide crypto policies only when you disable them through the workaround. ([BZ#1919155](#))

5.7.5. Networking

The iptables utility now requests module loading for commands that update a chain regardless of the NLM_F_CREATE flag

Previously, when setting a chain's policy, the **iptables-nft** utility generated a **NEWCHAIN** message but did not set the **NLM_F_CREATE** flag. As a consequence, the RHEL 8 kernel did not load any modules and the resulting update chain command failed if the associated kernel modules were not manually loaded. With this update, the **iptables-nft** utility now requests module loading for all commands that update a chain and users are able to set a chain's policy using the **iptables-nft** utility without manually loading the associated modules.

([BZ#1812666](#))

Support for updating packet/byte counters in the kernel was changed incorrectly between RHEL 7 and RHEL 8

When referring to an **ipset** command with enabled counters from an **iptables** rule, which specifies additional constraints on matching **ipset** entries, the **ipset** counters are updated only if all the additional constraints match. This is also problematic with **--packets-gt** or **--bytes-gt** constraints.

As a result, when migrating an **iptables** ruleset from RHEL 7 to RHEL 8, the rules involving **ipset** lookups may stop working and need to be adjusted. To work around this problem, avoid using the **--packets-gt** or **--bytes-gt** options and replace them with the **--packets-lt** or **--bytes-lt** options.

([BZ#1806882](#))

Unloading XDP programs fails on Netronome network cards that use the nfp driver

The **nfp** driver for Netronome network cards contains a bug. Therefore, unloading eXpress Data Path (XDP) programs fails if you use such cards and load the XDP program using the

IFLA_XDP_EXPECTED_FD feature with the **XDP_FLAGS_REPLACE** flag. For example, this bug affects XDP programs that are loaded using the **libxdp** library. Currently, there is no workaround available for the problem.

(BZ#1880268)

Anaconda does not have network access when using DHCP in the **ip** boot option

The initial RAM disk (**initrd**) uses NetworkManager to manage networking. The **dracut** NetworkManager module provided by the RHEL 8.3 ISO file incorrectly assumes that the first field of the **ip** option in the Anaconda boot options is always set. As a consequence, if you use DHCP and set **ip=::::<host_name>::dhcp**, NetworkManager does not retrieve an IP address, and the network is not available in Anaconda.

You have the following options to work around the problem:

1. Set the first field in the **ip`option to `.** (period):

```
ip=.::::<host_name>::dhcp
```

Note that this work around will not work in future versions of RHEL when the problem has been fixed.

2. Re-create the **boot.iso** file using the latest packages from the BaseOS repository that contains a fix for the bug: .

```
# lorax '--product=Red Hat Enterprise Linux' --version=8.3 --release=8.3 \
--source=<URL_to_BaseOS_repository> \
--source=<URL_to_AppStream_repository> \
--nomacboot --buildarch=x86_64 '--volid=RHEL 8.3' <output_directory>
```

. Note that Red Hat does not support self-created ISO files.

As a result, RHEL retrieves an IP address from the DHCP server, and network access is available in Anaconda.

(BZ#1902791)

5.7.6. Kernel

The **tboot-1.9.12-2** utility causes a boot failure in RHEL 8

The **tboot** utility of version 1.9.12-2 causes some systems with Trusted Platform Module (TPM) 2.0 to fail to boot in legacy mode. As a consequence, the system halts once it attempts to boot from the tboot Grand Unified Bootloader (GRUB) entry. To workaround this problem, downgrade to **tboot** of version 1.9.10.

(BZ#1947839)

The kernel returns false positive warnings on IBM Z systems

In RHEL 8, IBM Z systems are missing a whitelist entry for the **ZONE_DMA** memory zone to allow user access. Consequently, the kernel returns false positive warnings such as:

```
...
Bad or missing usercopy whitelist? Kernel memory exposure attempt detected from SLUB object
```

```
'dma-kmalloc-192' (offset 0, size 144)!  
WARNING: CPU: 0 PID: 8519 at mm/usercopy.c:83 usercopy_warn+0xac/0xd8  
...
```

The warnings appear when accessing certain system information through the **sysfs** interface. For example, by running the **debuginfo.sh** script.

To work around this problem, add the **hardened_usercopy=off** parameter to the kernel command line.

As a result, no warning messages are displayed in the described scenario.

(BZ#1660290)

The rngd service busy wait causes total CPU consumption in FIPS mode

A new kernel entropy source for FIPS mode has been added for kernels starting with version 4.18.0-193.10. Consequently, when in FIPS mode, the **rngd** service busy waits on the **poll()** system call for the **/dev/random** device, thereby causing consumption of 100% of CPU time. To work around this problem, stop and disable **rngd** by running:

```
# systemctl stop rngd  
# systemctl disable rngd
```

As a result, **rngd** no longer busy waits on **poll()** in the described scenario.

(BZ#1884857)

softirq changes can cause the localhost interface to drop UDP packets when under heavy load

Changes in the Linux kernel's software interrupt (**softirq**) handling are done to reduce denial of service (DOS) effects. Consequently, this leads to situations where the localhost interface drops User Datagram Protocol (UDP) packets under heavy load.

To work around this problem, increase the size of the network device backlog buffer to value 6000:

```
echo 6000 > /proc/sys/net/core/netdev_max_backlog
```

In Red Hat tests, this value was sufficient to prevent packet loss. More heavily loaded systems might require larger backlog values. Increased backlogs have the effect of potentially increasing latency on the localhost interface.

The result is to increase the buffer and allow more packets to be waiting for processing, which reduces the chances of dropping localhost packets.

(BZ#1779337)

A vmcore capture fails after memory hot-plug or unplug operation

After performing the memory hot-plug or hot-unplug operation, the event comes after updating the device tree which contains memory layout information. Thereby the **makedumpfile** utility tries to access a non-existent physical address. The problem appears if all of the following conditions meet:

- A little-endian variant of IBM Power System runs RHEL 8.
- The **kdump** or **fadump** service is enabled on the system.

Consequently, the capture kernel fails to save **vmcore** if a kernel crash is triggered after the memory hot-plug or hot-unplug operation.

To work around this problem, restart the **kdump** service after hot-plug or hot-unplug:

```
# systemctl restart kdump.service
```

As a result, **vmcore** is successfully saved in the described scenario.

(BZ#1793389)

Using **irqpoll** causes **vmcore** generation failure

Due to an existing problem with the **nvme** driver on the 64-bit ARM architectures that run on the Amazon Web Services (AWS) cloud platforms, the **vmcore** generation fails when you provide the **irqpoll** kernel command line parameter to the first kernel. Consequently, no **vmcore** file is dumped in the **/var/crash/** directory after a kernel crash. To work around this problem:

1. Add **irqpoll** to the **KDUMP_COMMANDLINE_REMOVE** key in the **/etc/sysconfig/kdump** file.
2. Restart the **kdump** service by running the **systemctl restart kdump** command.

As a result, the first kernel boots correctly and the **vmcore** file is expected to be captured upon the kernel crash.

Note that the **kdump** service can use a significant amount of crash kernel memory to dump the **vmcore** file. Ensure that the capture kernel has sufficient memory available for the **kdump** service.

(BZ#1654962)

Debug kernel fails to boot in crash capture environment in RHEL 8

Due to memory-demanding nature of the debug kernel, a problem occurs when the debug kernel is in use and a kernel panic is triggered. As a consequence, the debug kernel is not able to boot as the capture kernel, and a stack trace is generated instead. To work around this problem, increase the crash kernel memory accordingly. As a result, the debug kernel successfully boots in the crash capture environment.

(BZ#1659609)

zlib may slow down a **vmcore** capture in some compression functions

The **kdump** configuration file uses the **lzo** compression format (**makedumpfile -l**) by default. When you modify the configuration file using the **zlib** compression format, (**makedumpfile -c**) it is likely to bring a better compression factor at the expense of slowing down the **vmcore** capture process. As a consequence, it takes the **kdump** upto four times longer to capture a **vmcore** with **zlib**, as compared to **lzo**.

As a result, Red Hat recommends using the default **lzo** for cases where speed is the main driving factor. However, if the target machine is low on available space, **zlib** is a better option.

(BZ#1790635)

The HP NMI watchdog does not always generate a crash dump

In certain cases, the **hpwdt** driver for the HP NMI watchdog is not able to claim a non-maskable interrupt (NMI) generated by the HPE watchdog timer because the NMI was instead consumed by the **perfmon** driver.

The missing NMI is initiated by one of two conditions:

1. The **Generate NMI** button on the Integrated Lights-Out (iLO) server management software. This button is triggered by a user.
2. The **hpwdt** watchdog. The expiration by default sends an NMI to the server.

Both sequences typically occur when the system is unresponsive. Under normal circumstances, the NMI handler for both these situations calls the **kernel panic()** function and if configured, the **kdump** service generates a **vmcore** file.

Because of the missing NMI, however, **kernel panic()** is not called and **vmcore** is not collected.

In the first case (1.), if the system was unresponsive, it remains so. To work around this scenario, use the virtual **Power** button to reset or power cycle the server.

In the second case (2.), the missing NMI is followed 9 seconds later by a reset from the Automated System Recovery (ASR).

The HPE Gen9 Server line experiences this problem in single-digit percentages. The Gen10 at an even smaller frequency.

(BZ#1602962)

The **tuned-adm profile powersave** command causes the system to become unresponsive

Executing the **tuned-adm profile powersave** command leads to an unresponsive state of the Penguin Valkyrie 2000 2-socket systems with the older Thunderx (CN88xx) processors. Consequently, reboot the system to resume working. To work around this problem, avoid using the **powersave** profile if your system matches the mentioned specifications.

(BZ#1609288)

The default **7 4 1 7 printk** value sometimes causes temporary system unresponsiveness

The default **7 4 1 7 printk** value allows for better debugging of the kernel activity. However, when coupled with a serial console, this **printk** setting can cause intense I/O bursts that can lead to a RHEL system becoming temporarily unresponsive. To work around this problem, we have added a new **optimize-serial-console** TuneD profile, which reduces the default **printk** value to **4 4 1 7**. Users can instrument their system as follows:

```
# tuned-adm profile throughput-performance optimize-serial-console
```

Having a lower **printk** value persistent across a reboot reduces the likelihood of system hangs.

Note that this setting change comes at the expense of losing the extra debugging information.

For more information about the newly added feature, see [A new **optimize-serial-console** TuneD profile to reduce I/O to serial consoles by lowering the **printk** value.](#)

(JIRA:RHELPLAN-28940)

The kernel ACPI driver reports it has no access to a PCIe ECAM memory region

The Advanced Configuration and Power Interface (ACPI) table provided by firmware does not define a memory region on the PCI bus in the Current Resource Settings (**_CRS**) method for the PCI bus device. Consequently, the following warning message occurs during the system boot:

■


```
[ 2.817152] acpi PNP0A08:00: [Firmware Bug]: ECAM area [mem 0x30000000-0x31ffffff] not
reserved in ACPI namespace
[ 2.827911] acpi PNP0A08:00: ECAM at [mem 0x30000000-0x31ffffff] for [bus 00-1f]
```

However, the kernel is still able to access the **0x30000000-0x31ffffff** memory region, and can assign that memory region to the PCI Enhanced Configuration Access Mechanism (ECAM) properly. You can verify that PCI ECAM works correctly by accessing the PCIe configuration space over the 256 byte offset with the following output:

```
03:00.0 Non-Volatile memory controller: Sandisk Corp WD Black 2018/PC SN720 NVMe SSD (prog-
if 02 [NVM Express])
...
Capabilities: [900 v1] L1 PM Substates
    L1SubCap: PCI-PM_L1.2- PCI-PM_L1.1- ASPM_L1.2+ ASPM_L1.1- L1_PM_Substates+
        PortCommonModeRestoreTime=255us PortTPowerOnTime=10us
    L1SubCtl1: PCI-PM_L1.2- PCI-PM_L1.1- ASPM_L1.2- ASPM_L1.1-
        T_CommonMode=0us LTR1.2_Threshold=0ns
    L1SubCtl2: T_PwrOn=10us
```

As a result, you can ignore the warning message.

For more information about the problem, see the "[Firmware Bug: ECAM area **mem 0x30000000-0x31ffffff** not reserved in ACPI namespace](#)" appears during system boot" solution.

(BZ#1868526)

The OPEN MPI library may trigger run-time failures with default PML

In OPEN Message Passing Interface (OPEN MPI) implementation 4.0.x series, Unified Communication X (UCX) is the default point-to-point communicator (PML). The later versions of OPEN MPI 4.0.x series deprecated **openib** Byte Transfer Layer (BTL).

However, OPEN MPI, when run over a **homogeneous** cluster (same hardware and software configuration), UCX still uses **openib** BTL for MPI one-sided operations. As a consequence, this may trigger execution errors. To work around this problem:

- Run the **mpirun** command using following parameters:

```
-mca btl openib -mca pml ucx -x UCX_NET_DEVICES=mlx5_ib0
```

where,

- The **-mca btl openib** parameter disables **openib** BTL
- The **-mca pml ucx** parameter configures OPEN MPI to use **ucx** PML.
- The **x UCX_NET_DEVICES=** parameter restricts UCX to use the specified devices

The OPEN MPI, when run over a **heterogeneous** cluster (different hardware and software configuration), it uses UCX as the default PML. As a consequence, this may cause the OPEN MPI jobs to run with erratic performance, unresponsive behavior, or crash failures. To work around this problem, set the UCX priority as:

- Run the **mpirun** command using following parameters:

```
-mca pml_ucx_priority 5
```

As a result, the OPEN MPI library is able to choose an alternative available transport layer over UCX.

(BZ#1866402)

5.7.7. File systems and storage

The **/boot** file system cannot be placed on LVM

You cannot place the **/boot** file system on an LVM logical volume. This limitation exists for the following reasons:

- On EFI systems, the *EFI System Partition* conventionally serves as the **/boot** file system. The uEFI standard requires a specific GPT partition type and a specific file system type for this partition.
- RHEL 8 uses the *Boot Loader Specification* (BLS) for system boot entries. This specification requires that the **/boot** file system is readable by the platform firmware. On EFI systems, the platform firmware can read only the **/boot** configuration defined by the uEFI standard.
- The support for LVM logical volumes in the GRUB 2 boot loader is incomplete. Red Hat does not plan to improve the support because the number of use cases for the feature is decreasing due to standards such as uEFI and BLS.

Red Hat does not plan to support **/boot** on LVM. Instead, Red Hat provides tools for managing system snapshots and rollback that do not need the **/boot** file system to be placed on an LVM logical volume.

(BZ#1496229)

LVM no longer allows creating volume groups with mixed block sizes

LVM utilities such as **vgcreate** or **vgextend** no longer allow you to create volume groups (VGs) where the physical volumes (PVs) have different logical block sizes. LVM has adopted this change because file systems fail to mount if you extend the underlying logical volume (LV) with a PV of a different block size.

To re-enable creating VGs with mixed block sizes, set the **allow_mixed_block_sizes=1** option in the **lvm.conf** file.

(BZ#1768536)

Limitations of LVM **writecache**

The **writecache** LVM caching method has the following limitations, which are not present in the **cache** method:

- You cannot name a **writecache** logical volume when using **pvmove** commands.
- You cannot use logical volumes with **writecache** in combination with thin pools or VDO.

The following limitation also applies to the **cache** method:

- You cannot resize a logical volume while **cache** or **writecache** is attached to it.

(JIRA:RHELPLAN-27987, BZ#1798631, BZ#1808012)

LVM mirror devices that store a LUKS volume sometimes become unresponsive

Mirrored LVM devices with a segment type of **mirror** that store a LUKS volume might become unresponsive under certain conditions. The unresponsive devices reject all I/O operations.

To work around the issue, Red Hat recommends that you use LVM RAID 1 devices with a segment type of **raid1** instead of **mirror** if you need to stack LUKS volumes on top of resilient software-defined storage.

The **raid1** segment type is the default RAID configuration type and replaces **mirror** as the recommended solution.

To convert **mirror** devices to **raid**, see [Converting a mirrored LVM device to a RAID1 logical volume](#) .

(BZ#1730502)

An NFS 4.0 patch can result in reduced performance under an open-heavy workload

Previously, a bug was fixed that, in some cases, could cause an NFS open operation to overlook the fact that a file had been removed or renamed on the server. However, the fix may cause slower performance with workloads that require many open operations. To work around this problem, it might help to use NFS version 4.1 or higher, which have been improved to grant delegations to clients in more cases, allowing clients to perform open operations locally, quickly, and safely.

(BZ#1748451)

5.7.8. Dynamic programming languages, web and database servers

getpwnam() might fail when called by a 32-bit application

When a user of NIS uses a 32-bit application that calls the **getpwnam()** function, the call fails if the **nss_nis.i686** package is missing. To work around this problem, manually install the missing package by using the **yum install nss_nis.i686** command.

(BZ#1803161)

Symbol conflicts between OpenLDAP libraries might cause crashes in **httpd**

When both the **libldap** and **libldap_r** libraries provided by OpenLDAP are loaded and used within a single process, symbol conflicts between these libraries might occur. Consequently, Apache **httpd** child processes using the PHP **ldap** extension might terminate unexpectedly if the **mod_security** or **mod_auth_openidc** modules are also loaded by the **httpd** configuration.

With this update to the Apache Portable Runtime (APR) library, you can work around the problem by setting the **APR_DEEPCBIND** environment variable, which enables the use of the **RTLD_DEEPCBIND** dynamic linker option when loading **httpd** modules. When the **APR_DEEPCBIND** environment variable is enabled, crashes no longer occur in **httpd** configurations that load conflicting libraries.

(BZ#1819607)

PAM plug-in does not work in MariaDB

MariaDB 10.3 provides the Pluggable Authentication Modules (PAM) plug-in version 1.0. The **MariaDB** PAM plug-in version 1.0 does not work in RHEL 8. To work around this problem, use the PAM plug-in version 2.0 provided by the **mariadb:10.5** module stream, which is available with RHEL 8.4.

(BZ#1942330)

5.7.9. Identity Management

Installing KRA fails if all KRA members are hidden replicas

The **ipa-kra-install** utility fails on a cluster where the Key Recovery Authority (KRA) is already present, if the first KRA instance is installed on a hidden replica. Consequently, you cannot add further KRA instances to the cluster.

To work around this problem, unhide the hidden replica that has the KRA role before you add new KRA instances. You can hide it again when **ipa-kra-install** completes successfully.

([BZ#1816784](#))

Using the **cert-fix** utility with the **--agent-uid pkidbuser** option breaks Certificate System

Using the **cert-fix** utility with the **--agent-uid pkidbuser** option corrupts the LDAP configuration of Certificate System. As a consequence, Certificate System might become unstable and manual steps are required to recover the system.

([BZ#1729215](#))

Certificates issued by PKI ACME Responder connected to PKI CA may fail OCSP validation

The default ACME certificate profile provided by PKI CA contains a sample OCSP URL that does not point to an actual OCSP service. As a consequence, if PKI ACME Responder is configured to use a PKI CA issuer, the certificates issued by the responder may fail OCSP validation.

To work around this problem, you need to set the **policyset.serverCertSet.5.default.params.authInfoAccessADLocation_0** property to a blank value in the **/usr/share/pki/ca/profiles/ca/acmeServerCert.cfg** configuration file:

1. In the ACME Responder configuration file, change the line
policyset.serverCertSet.5.default.params.authInfoAccessADLocation_0=http://ocsp.example.com to **policyset.serverCertSet.5.default.params.authInfoAccessADLocation_0=.**
2. Restart the service and regenerate the certificate.

As a result, PKI CA will generate ACME certificates with an autogenerated OCSP URL that points to an actual OCSP service.

([BZ#1868233](#))

FreeRADIUS silently truncates Tunnel-Passwords longer than 249 characters

If a Tunnel-Password is longer than 249 characters, the FreeRADIUS service silently truncates it. This may lead to unexpected password incompatibilities with other systems.

To work around the problem, choose a password that is 249 characters or fewer.

([BZ#1723362](#))

The **/var/log/lastlog** sparse file on IdM hosts can cause performance problems

During the IdM installation, a range of 200,000 UIDs from a total of 10,000 possible ranges is randomly selected and assigned. Selecting a random range in this way significantly reduces the probability of conflicting IDs in case you decide to merge two separate IdM domains in the future.

However, having high UIDs can create problems with the **/var/log/lastlog** file. For example, if a user with the UID of 1280000008 logs in to an IdM client, the local **/var/log/lastlog** file size increases to almost 400 GB. Although the actual file is sparse and does not use all that space, certain applications are not designed to identify sparse files by default and may require a specific option to handle them. For example, if the setup is complex and a backup and copy application does not handle sparse files correctly, the file is copied as if its size was 400 GB. This behavior can cause performance problems.

To work around this problem:

- In case of a standard package, refer to its documentation to identify the option that handles sparse files.
- In case of a custom application, ensure that it is able to manage sparse files such as **/var/log/lastlog** correctly.

(JIRA:RHELPLAN-59111)

Potential risk when using the default value for **ldap_id_use_start_tls** option

When using **ldap://** without TLS for identity lookups, it can pose a risk for an attack vector. Particularly a man-in-the-middle (MITM) attack which could allow an attacker to impersonate a user by altering, for example, the UID or GID of an object returned in an LDAP search.

Currently, the SSSD configuration option to enforce TLS, **ldap_id_use_start_tls**, defaults to **false**. Ensure that your setup operates in a trusted environment and decide if it is safe to use unencrypted communication for **id_provider = ldap**. Note **id_provider = ad** and **id_provider = ipa** are not affected as they use encrypted connections protected by SASL and GSSAPI.

If it is not safe to use unencrypted communication, enforce TLS by setting the **ldap_id_use_start_tls** option to **true** in the **/etc/sss/sss.conf** file. The default behavior is planned to be changed in a future release of RHEL.

(JIRA:RHELPLAN-155168)

5.7.10. Desktop

Disabling **flatpak** repositories from Software Repositories is not possible

Currently, it is not possible to disable or remove **flatpak** repositories in the Software Repositories tool in the GNOME Software utility.

([BZ#1668760](#))

Drag-and-drop does not work between desktop and applications

Due to a bug in the **gnome-shell-extensions** package, the drag-and-drop functionality does not currently work between desktop and applications. Support for this feature will be added back in a future release.

([BZ#1717947](#))

Generation 2 RHEL 8 virtual machines sometimes fail to boot on Hyper-V Server 2016 hosts

When using RHEL 8 as the guest operating system on a virtual machine (VM) running on a Microsoft Hyper-V Server 2016 host, the VM in some cases fails to boot and returns to the GRUB boot menu. In addition, the following error is logged in the Hyper-V event log:

The guest operating system reported that it failed with the following error code: 0x1E

This error occurs due to a UEFI firmware bug on the Hyper-V host. To work around this problem, use Hyper-V Server 2019 as the host.

([BZ#1583445](#))

5.7.11. Graphics infrastructures

radeon fails to reset hardware correctly

The **radeon** kernel driver currently does not reset hardware in the **kexec** context correctly. Instead, **radeon** falls over, which causes the rest of the **kdump** service to fail.

To work around this problem, disable **radeon** in **kdump** by adding the following line to the **/etc/kdump.conf** file:

```
dracut_args --omit-drivers "radeon"  
force_rebuild 1
```

Restart the machine and **kdump**. After starting **kdump**, the **force_rebuild 1** line may be removed from the configuration file.

Note that in this scenario, no graphics will be available during **kdump**, but **kdump** will work successfully.

(BZ#1694705)

Multiple HDR displays on a single MST topology may not power on

On systems using NVIDIA Turing GPUs with the **nouveau** driver, using a **DisplayPort** hub (such as a laptop dock) with multiple monitors which support HDR plugged into it may result in failure to turn on. This is due to the system erroneously thinking there is not enough bandwidth on the hub to support all of the displays.

(BZ#1812577)

Unable to run graphical applications using `sudo` command

When trying to run graphical applications as a user with elevated privileges, the application fails to open with an error message. The failure happens because **Xwayland** is restricted by the **Xauthority** file to use regular user credentials for authentication.

To work around this problem, use the **sudo -E** command to run graphical applications as a **root** user.

(BZ#1673073)

VNC Viewer displays wrong colors with the 16-bit color depth on IBM Z

The VNC Viewer application displays wrong colors when you connect to a VNC session on an IBM Z server with the 16-bit color depth.

To work around the problem, set the 24-bit color depth on the VNC server. With the **Xvnc** server, replace the **-depth 16** option with **-depth 24** in the **Xvnc** configuration.

As a result, VNC clients display the correct colors but use more network bandwidth with the server.

(BZ#1886147)

Hardware acceleration is not supported on ARM

Built-in graphics drivers do not support hardware acceleration or the Vulkan API on the 64-bit ARM architecture.

To enable hardware acceleration or Vulkan on ARM, install the proprietary Nvidia driver.

(JIRA:RHELPLAN-57914)

The RHEL installer becomes unresponsive with NVIDIA Ampere

RHEL 8.3.0 does not support the NVIDIA Ampere GPUs. If you start the RHEL installation on a system that has an NVIDIA Ampere GPU, the installer becomes unresponsive. As a consequence, the installation cannot finish successfully.

The NVIDIA Ampere family includes the following GPU models:

- GeForce RTX 3060 Ti
- GeForce RTX 3070
- GeForce RTX 3080
- GeForce RTX 3090
- RTX A6000
- NVIDIA A40
- NVIDIA A100
- NVIDIA A100 80GB

To work around the problem, disable the **nouveau** graphics driver and install RHEL in text mode:

1. Boot into the boot menu of the installer.
2. Add the **nouveau.modeset=0** option on the kernel command line.
For details, see [Editing boot options](#).
3. Install RHEL on the system.
4. Boot into the newly installed RHEL. At the boot menu, add the **nouveau.modeset=0** option on the kernel command line.
5. Disable the **nouveau** driver permanently:

```
# echo 'blacklist nouveau' >> /etc/modprobe.d/blacklist.conf
```

As a result, the installation has finished successfully and RHEL now runs in text mode.

Optionally, you can install the proprietary NVIDIA GPU driver to enable graphics. For instructions, see [How to install the NVIDIA proprietary driver on RHEL 8](#).

(BZ#1903890)

5.7.12. The web console

Unprivileged users can access the Subscriptions page

If a non-administrator navigates to the **Subscriptions** page of the web console, the web console displays a generic error message **Cockpit had an unexpected internal error**.

To work around this problem, sign in to the web console with a privileged user and make sure to check the **Reuse my password for privileged tasks** checkbox.

([BZ#1674337](#))

5.7.13. Red Hat Enterprise Linux System Roles

oVirt input and the elasticsearch output functionalities are not supported in System Roles Logging

The **oVirt** input and the **elasticsearch** output are not supported in System Roles Logging although they are mentioned in the README file. There is no workaround available at the moment.

([BZ#1889468](#))

5.7.14. Virtualization

Displaying multiple monitors of virtual machines that use Wayland is not possible with QXL

Using the **remote-viewer** utility to display more than one monitor of a virtual machine (VM) that is using the Wayland display server causes the VM to become unresponsive and the *Waiting for display* status message to be displayed indefinitely.

To work around this problem, use **virtio-gpu** instead of **qxl** as the GPU device for VMs that use Wayland.

([BZ#1642887](#))

virsh iface-* commands do not work consistently

Currently, **virsh iface-*** commands, such as **virsh iface-start** and **virsh iface-destroy**, frequently fail due to configuration dependencies. Therefore, it is recommended not to use **virsh iface-*** commands for configuring and managing host network connections. Instead, use the NetworkManager program and its related management applications.

([BZ#1664592](#))

Virtual machines sometimes fail to start when using many virtio-blk disks

Adding a large number of virtio-blk devices to a virtual machine (VM) may exhaust the number of interrupt vectors available in the platform. If this occurs, the VM's guest OS fails to boot, and displays a **dracut-initqueue[392]: Warning: Could not boot** error.

([BZ#1719687](#))

Attaching LUN devices to virtual machines using virtio-blk does not work

The q35 machine type does not support transitional virtio 1.0 devices, and RHEL 8 therefore lacks support for features that were deprecated in virtio 1.0. In particular, it is not possible on a RHEL 8 host to send SCSI commands from virtio-blk devices. As a consequence, attaching a physical disk as a LUN device to a virtual machine fails when using the virtio-blk controller.

Note that physical disks can still be passed through to the guest operating system, but they should be configured with the **device='disk'** option rather than **device='lun'**.

([BZ#1777138](#))

Virtual machines using Cooperlake cannot boot when TSX is disabled on the host

Virtual machines (VMs) that use the **Cooperlake** CPU model currently fail to boot when the **TSX** CPU flag is disabled on the host. Instead, the host displays the following error message:

```
the CPU is incompatible with host CPU: Host CPU does not provide required features: hle, rtm
```

To make VMs with **Cooperlake** usable on such host, disable the HLE, RTM, and TAA_NO flags in the VM configuration in the VM's XML configuration:

```
<feature policy='disable' name='hle'/>
<feature policy='disable' name='rtm'/>
<feature policy='disable' name='taa-no'/>
```

([BZ#1860743](#))

Virtual machines sometimes cannot boot on Witherspoon hosts

Virtual machines (VMs) that use the **pseries-rhel7.6.0-sxxm** machine type in some cases fail to boot on *Power9 S922LC for HPC* hosts (also known as Witherspoon) that use the DD2.2 or DD2.3 CPU.

Attempting to boot such a VM instead generates the following error message:

```
qemu-kvm: Requested safe indirect branch capability level not supported by kvm
```

To work around this problem, configure the VM's XML configuration as follows:

```
<domain type='qemu' xmlns:qemu='http://libvirt.org/schemas/domain/qemu/1.0'>
  <qemu:commandline>
    <qemu:arg value='-machine'/>
    <qemu:arg value='cap-ibs=workaround'/>
  </qemu:commandline>
```

([BZ#1732726](#))

5.7.15. RHEL in cloud environments

GPU problems on Azure NV6 instances

When running RHEL 8 as a guest operating system on a Microsoft Azure NV6 instance, resuming the virtual machine (VM) from hibernation sometimes causes the VM's GPU to work incorrectly. When this occurs, the kernel logs the following message:

```
hv_irq_unmask() failed: 0x5
```

([BZ#1846838](#))

kdump sometimes does not start on Azure and Hyper-V

On RHEL 8 guest operating systems hosted on the Microsoft Azure or Hyper-V hypervisors, starting the **kdump** kernel in some cases fails when post-exec notifiers are enabled.

To work around this problem, disable crash kexec post notifiers:

```
# echo N > /sys/module/kernel/parameters/crash_kexec_post_notifiers
```

(BZ#1865745)

Setting static IP in a RHEL 8 virtual machine on a VMWare host does not work

Currently, when using RHEL 8 as a guest operating system of a virtual machine (VM) on a VMWare host, the DatasourceOVF function does not work correctly. As a consequence, if you use the **cloud-init** utility to set the VM's network to static IP and then reboot the VM, the VM's network will be changed to DHCP.

(BZ#1750862)

Core dumping RHEL 8 virtual machines with certain NICs to a remote machine on Azure takes longer than expected

Currently, using the **kdump** utility to save the core dump file of a RHEL 8 virtual machine (VM) on a Microsoft Azure hypervisor to a remote machine does not work correctly when the VM is using a NIC with enabled accelerated networking. As a consequence, the dump file is saved after approximately 200 seconds, instead of immediately. In addition, the following error message is logged on the console before the dump file is saved.

```
device (eth0): linklocal6: DAD failed for an EUI-64 address
```

(BZ#1854037)

TX/RX packet counters do not increase after virtual machines resume from hibernation

The **TX/RX** packet counters stop increasing when a RHEL 8 virtual machine (VM), with a CX4 VF NIC, resumes from hibernation on Microsoft Azure. To keep the counters working, restart the VM. Note that, doing so will reset the counters.

(BZ#1876527)

RHEL 8 virtual machines fail to resume from hibernation on Azure

The GUID of the virtual function (VF), **vmbus device**, changes when a RHEL 8 virtual machine (VM), with **SR-IOV** enabled, is hibernated and deallocated on Microsoft Azure. As a result, when the VM is restarted, it fails to resume and crashes. As a workaround, hard reset the VM using the Azure serial console.

(BZ#1876519)

Migrating a POWER9 guest from a RHEL 7-ALT host to RHEL 8 fails

Currently, migrating a POWER9 virtual machine from a RHEL 7-ALT host system to RHEL 8 becomes unresponsive with a "Migration status: active" status.

To work around this problem, disable Transparent Huge Pages (THP) on the RHEL 7-ALT host, which enables the migration to complete successfully.

(BZ#1741436)

5.7.16. Supportability

redhat-support-tool does not work with the FUTURE crypto policy

Because a cryptographic key used by a certificate on the Customer Portal API does not meet the requirements by the **FUTURE** system-wide cryptographic policy, the **redhat-support-tool** utility does not work with this policy level at the moment.

To work around this problem, use the **DEFAULT** crypto policy while connecting to the Customer Portal API.

([BZ#1802026](#))

5.7.17. Containers

UDICA is not expected to work with 1.0 stable stream

UDICA, the tool to generate SELinux policies for containers, is not expected to work with containers that are run via podman 1.0.x in the **container-tools:1.0** module stream.

(JIRA:RHELPLAN-25571)

podman system connection add does not automatically set the default connection

The **podman system connection add** command does not automatically set the first connection to be the default connection. To set the default connection, you must manually run the command **podman system connection default <connection_name>**.

([BZ#1881894](#))

CHAPTER 6. INTERNATIONALIZATION

6.1. RED HAT ENTERPRISE LINUX 8 INTERNATIONAL LANGUAGES

Red Hat Enterprise Linux 8 supports the installation of multiple languages and the changing of languages based on your requirements.

- East Asian Languages - Japanese, Korean, Simplified Chinese, and Traditional Chinese.
- European Languages - English, German, Spanish, French, Italian, Portuguese, and Russian.

The following table lists the fonts and input methods provided for various major languages.

Language	Default Font (Font Package)	Input Methods
English	dejavu-sans-fonts	
French	dejavu-sans-fonts	
German	dejavu-sans-fonts	
Italian	dejavu-sans-fonts	
Russian	dejavu-sans-fonts	
Spanish	dejavu-sans-fonts	
Portuguese	dejavu-sans-fonts	
Simplified Chinese	google-noto-sans-cjk-ttc-fonts, google-noto-serif-cjk-ttc-fonts	ibus-libpinyin, libpinyin
Traditional Chinese	google-noto-sans-cjk-ttc-fonts, google-noto-serif-cjk-ttc-fonts	ibus-libzhuyin, libzhuyin
Japanese	google-noto-sans-cjk-ttc-fonts, google-noto-serif-cjk-ttc-fonts	ibus-kkc, libkkc
Korean	google-noto-sans-cjk-ttc-fonts, google-noto-serif-cjk-ttc-fonts	ibus-hangul, libhangul

6.2. NOTABLE CHANGES TO INTERNATIONALIZATION IN RHEL 8

RHEL 8 introduces the following changes to internationalization compared to RHEL 7:

- Support for the **Unicode 11** computing industry standard has been added.
- Internationalization is distributed in multiple packages, which allows for smaller footprint installations. For more information, see [Using langpacks](#).

- A number of **glibc** locales have been synchronized with Unicode Common Locale Data Repository (CLDR).

APPENDIX A. LIST OF TICKETS BY COMPONENT

Bugzilla and JIRA IDs are listed in this document for reference. Bugzilla bugs that are publicly accessible include a link to the ticket.

Component	Tickets
389-ds-base	BZ#1816862 , BZ#1638875 , BZ#1728943
NetworkManager	BZ#1814746 , BZ#1626348
anaconda	BZ#1665428 , BZ#1775975 , BZ#1630299 , BZ#1823578 , BZ#1672405 , BZ#1644662 , BZ#1745064 , BZ#1821192 , BZ#1822880 , BZ#1862116 , BZ#1890261 , BZ#1891827 , BZ#1691319 , BZ#1931069
apr	BZ#1819607
authselect	BZ#1654018
bcc	BZ#1837906
bind	BZ#1818785
buildah-container	BZ#1627898
buildah	BZ#1806044
clevis	BZ#1716040 , BZ#1818780 , BZ#1436735 , BZ#1819767
cloud-init	BZ#1750862
cloud-utils-growpart	BZ#1846246
cockpit-session-recording	BZ#1826516
cockpit	BZ#1710731 , BZ#1666722
corosync-qdevice	BZ#1784200
crun	BZ#1841438
crypto-policies	BZ#1832743 , BZ#1660839
cyrus-sasl	BZ#1817054
distribution	BZ#1815402 , BZ#1657927

Component	Tickets
dnf	BZ#1793298 , BZ#1832869 , BZ#1842285
elfutils	BZ#1804321
fapolicyd	BZ#1897090 , BZ#1817413 , BZ#1714529
fence-agents	BZ#1830776 , BZ#1775847
firewalld	BZ#1790948 , BZ#1682913 , BZ#1809225 , BZ#1817205 , BZ#1809636
freeradius	BZ#1672285 , BZ#1859527 , BZ#1723362
gcc-toolset-10-gdb	BZ#1838777
gcc	BZ#1784758
gdb	BZ#1659535
git	BZ#1825114
glibc	BZ#1812756 , BZ#1743445 , BZ#1783303 , BZ#1642150 , BZ#1810146 , BZ#1748197 , BZ#1774115 , BZ#1807824 , BZ#1757354 , BZ#1836867 , BZ#1780204 , BZ#1821531 , BZ#1784525
gnome-session	BZ#1739556
gnome-shell-extensions	BZ#1717947
gnome-shell	BZ#1724302
gnome-software	BZ#1668760
gnutls	BZ#1677754 , BZ#1789392 , BZ#1849079 , BZ#1855803
go-toolset	BZ#1820596
gpgme	BZ#1829822
grafana-container	BZ#1823834
grafana-pcp	BZ#1807099
grafana	BZ#1807323
grub2	BZ#1583445

Component	Tickets
httpd	BZ#1209162
initial-setup	BZ#1676439
ipa-healthcheck	BZ#1852244
ipa	BZ#1816784 , BZ#1810154 , BZ#913799 , BZ#1651577 , BZ#1851139 , BZ#1664719 , BZ#1664718
iperf3	BZ#1665142 , BZ#1700497
jss	BZ#1821851
kernel-rt	BZ#1818138
kernel	BZ#1758323 , BZ#1812666 , BZ#1793389 , BZ#1694705 , BZ#1748451 , BZ#1654962 , BZ#1792125 , BZ#1708456 , BZ#1812577 , BZ#1757933 , BZ#1847837 , BZ#1791664 , BZ#1666538 , BZ#1602962 , BZ#1609288 , BZ#1730502 , BZ#1806882 , BZ#1660290 , BZ#1846838 , BZ#1865745 , BZ#1868526 , BZ#1884857 , BZ#1854037 , BZ#1876527 , BZ#1876519 , BZ#1823764 , BZ#1822085 , BZ#1735611 , BZ#1281843 , BZ#1828642 , BZ#1825414 , BZ#1761928 , BZ#1791041 , BZ#1796565 , BZ#1834769 , BZ#1785660 , BZ#1683394 , BZ#1817752 , BZ#1782831 , BZ#1821646 , BZ#1519039 , BZ#1627455 , BZ#1501618 , BZ#1495358 , BZ#1633143 , BZ#1503672 , BZ#1570255 , BZ#1696451 , BZ#1348508 , BZ#1778762 , BZ#1839311 , BZ#1783396 , BZ#1665295 , BZ#1658840 , BZ#1660627 , BZ#1569610
krb5	BZ#1791062 , BZ#1784655 , BZ#1820311 , BZ#1802334 , BZ#1877991
libbpf	BZ#1759154
libcap	BZ#1487388
libdb	BZ#1670768
libffi	BZ#1723951
libgnome-keyring	BZ#1607766
libkcapi	BZ#1683123
libmaxminddb	BZ#1642001
libpcap	BZ#1806422

Component	Tickets
libreswan	BZ#1544463 , BZ#1820206
libseccomp	BZ#1770693
libselinux-python-2.8-module	BZ#1666328
libssh	BZ#1804797
libvirt	BZ#1664592 , BZ#1528684
lldb	BZ#1841073
llvm-toolset	BZ#1820587
llvm	BZ#1820319
lshw	BZ#1794049
lvm2	BZ#1496229 , BZ#1768536 , BZ#1598199 , BZ#1541165 , JIRA:RHELPLAN-39320
mariadb	BZ#1942330
memcached	BZ#1809536
mesa	BZ#1886147
microdnf	BZ#1781126
mod_http2	BZ#1814236
nfs-utils	BZ#1817756 , BZ#1592011
nginx	BZ#1668717 , BZ#1826632
nmstate	BZ#1674456
nss_nis	BZ#1803161
nss	BZ#1817533 , BZ#1645153
opencryptoki	BZ#1780293

Component	Tickets
openmpi	BZ#1866402
opensc	BZ#1810660
openscap	BZ#1803116 , BZ#1870087 , BZ#1795563 , BZ#1824152 , BZ#1829761
openssh	BZ#1744108
openssl	BZ#1685470, BZ#1810911
oscap-anaconda-addon	BZ#1816199, BZ#1665082, BZ#1674001 , BZ#1691305, BZ#1787156 , BZ#1843932 , BZ#1834716
pacemaker	BZ#1828488 , BZ#1784601 , BZ#1837747, BZ#1718324
papi	BZ#1807346, BZ#1664056 , BZ#1726070
pcp-container	BZ#1497296
pcp	BZ#1792971
pcs	BZ#1817547, BZ#1684676 , BZ#1839637 , BZ#1619620
perl-5.30-module	BZ#1713592
perl-IO-Socket-SSL	BZ#1824222
perl-libwww-perl	BZ#1781177
php	BZ#1797661
pki-core	BZ#1729215 , BZ#1868233 , BZ#1770322, BZ#1824948
podman	BZ#1804193 , BZ#1881894 , BZ#1627899
powertop	BZ#1783110
pykickstart	BZ#1637872
python38	BZ#1847416
qemu-kvm	BZ#1719687 , BZ#1860743 , JIRA:RHELPLAN-45901, BZ#1651994
rear	BZ#1843809 , BZ#1729502 , BZ#1743303

Component	Tickets
redhat-support-tool	BZ#1802026
resource-agents	BZ#1814896
rhel-system-roles-sap	BZ#1844190 , BZ#1660832
rhel-system-roles	BZ#1889468 , BZ#1822158 , BZ#1677739
rpm	BZ#1688849
rsyslog	BZ#1659383 , JIRA:RHELPLAN-10431 , BZ#1679512 , BZ#1713427
ruby-2.7-module	BZ#1817135
ruby	BZ#1846113
rust-toolset	BZ#1820593
samba	BZ#1817557 , JIRA:RHELPLAN-13195
scap-security-guide	BZ#1843913 , BZ#1858866 , BZ#1750755 , BZ#1760734 , BZ#1832760 , BZ#1815007
scap-workbench	BZ#1640715
selinux-policy	BZ#1826788 , BZ#1746398 , BZ#1776873 , BZ#1772852 , BZ#1641631 , BZ#1860443
setools	BZ#1820079
skopeo-container	BZ#1627900
smartmontools	BZ#1671154
spice	BZ#1849563
squid	BZ#1829467
sssd	BZ#1827615 , BZ#1793727
stratis-cli	BZ#1734496
stunnel	BZ#1808365
subscription-manager	BZ#1674337

Component	Tickets
sudo	BZ#1786990
systemtap	BZ#1804319
tang	BZ#1716039
tcpdump	BZ#1804063
tigervnc	BZ#1806992
tpm2-tools	BZ#1789682
tuned	BZ#1792264 , BZ#1840689 , BZ#1746957
udica	BZ#1763210
usbguard	BZ#1738590 , BZ#1667395 , BZ#1683567
valgrind	BZ#1804324
wayland	BZ#1673073
xdp-tools	BZ#1880268 , BZ#1820670
xorg-x11-drv-qxl	BZ#1642887
xorg-x11-server	BZ#1698565
yum	BZ#1788154

Component	Tickets
other	<p>JIRA:RHELPLAN-45950, JIRA:RHELPLAN-57572, BZ#1640697, BZ#1659609, BZ#1687900, BZ#1697896, BZ#1790635, BZ#1823398, BZ#1757877, JIRA:RHELPLAN-25571, BZ#1777138, JIRA:RHELPLAN-27987, JIRA:RHELPLAN-28940, JIRA:RHELPLAN-34199, JIRA:RHELPLAN-57914, BZ#1897383, BZ#1900019, BZ#1839151, BZ#1780124, JIRA:RHELPLAN-42395, BZ#1889736, BZ#1842656, JIRA:RHELPLAN-45959, JIRA:RHELPLAN-45958, JIRA:RHELPLAN-45957, JIRA:RHELPLAN-45956, JIRA:RHELPLAN-45952, JIRA:RHELPLAN-45945, JIRA:RHELPLAN-45939, JIRA:RHELPLAN-45937, JIRA:RHELPLAN-45936, JIRA:RHELPLAN-45930, JIRA:RHELPLAN-45926, JIRA:RHELPLAN-45922, JIRA:RHELPLAN-45920, JIRA:RHELPLAN-45918, JIRA:RHELPLAN-45916, JIRA:RHELPLAN-45915, JIRA:RHELPLAN-45911, JIRA:RHELPLAN-45910, JIRA:RHELPLAN-45909, JIRA:RHELPLAN-45908, JIRA:RHELPLAN-45906, JIRA:RHELPLAN-45904, JIRA:RHELPLAN-45900, JIRA:RHELPLAN-45899, JIRA:RHELPLAN-45884, JIRA:RHELPLAN-37573, JIRA:RHELPLAN-37570, JIRA:RHELPLAN-49954, JIRA:RHELPLAN-50002, JIRA:RHELPLAN-43531, JIRA:RHELPLAN-48838, BZ#1873567, BZ#1866695, JIRA:RHELPLAN-14068, JIRA:RHELPLAN-7788, JIRA:RHELPLAN-40469, JIRA:RHELPLAN-42617, JIRA:RHELPLAN-30878, JIRA:RHELPLAN-37517, JIRA:RHELPLAN-55009, JIRA:RHELPLAN-42396, BZ#1836211, JIRA:RHELPLAN-57564, JIRA:RHELPLAN-57567, BZ#1890499, JIRA:RHELPLAN-40234, JIRA:RHELPLAN-56676, JIRA:RHELPLAN-14754, JIRA:RHELPLAN-51289, BZ#1893174, BZ#1690207, JIRA:RHELPLAN-1212, BZ#1559616, BZ#1889737, BZ#1812552, JIRA:RHELPLAN-14047, BZ#1769727, JIRA:RHELPLAN-27394, JIRA:RHELPLAN-27737, JIRA:RHELPLAN-41549, BZ#1642765, JIRA:RHELPLAN-10304, BZ#1646541, BZ#1647725, BZ#1686057, BZ#1748980, BZ#1827628, BZ#1871025, BZ#1871953, BZ#1874892, BZ#1893767, JIRA:RHELPLAN-60226</p>

APPENDIX B. REVISION HISTORY

0.3-5

Thu Apr 27 2023, Gabriela Fialová (gfialova@redhat.com)

- Added a known issue [JIRA:RHELPLAN-155168](#) (Identity Management).

0.3-4

Fri Apr 29 2022, Lenka Špačková (lspackova@redhat.com)

- Updated [Deprecated functionality](#) introduction.
- Fixed typo in [BZ#1605216](#).
- Fixed broken links.

0.3-3

Tue Apr 05 2022, Jaroslav Klech (jklech@redhat.com)

- Added a bug fix [BZ#1666538](#) (Kernel).

0.3-2

Thu Mar 17 2022, Jaroslav Klech (jklech@redhat.com)

- Added a known issue (Kernel).

0.3-1

Fri Feb 04 2022, Lucie Maňásková ([lmanasko@redhat.com](mailto:لماناسكو@redhat.com))

- Added deprecated functionality [BZ#1794513](#) (Filesystems and storage).

0.3-0

Thu Dec 23 2021, Lenka Špačková (lspackova@redhat.com)

- Added information about the Soft-RoCE driver, **rdma_rxe**, to Technology Previews [BZ#1605216](#) and Deprecated Functionality [BZ#1878207](#) (Kernel).

0.2-9

Thu Oct 07 2021, Lenka Špačková (lspackova@redhat.com)

- Updated the known issue [BZ#1942330](#) (Dynamic programming languages, web and database servers).

0.2-8

Tue Oct 05 2021, Lucie Maňásková ([lmanasko@redhat.com](mailto:لماناسكو@redhat.com))

- Added a deprecated functionality [BZ#1999620](#) (Shells and command-line tools).

0.2-7

Thu Aug 19 2021, Lucie Maňásková ([lmanasko@redhat.com](mailto:لماناسكو@redhat.com))

- Added the [Package management with YUM/DNF](#) to the Distribution chapter.

- Updated the text of [BZ#1708456](#) (Krenel).

0.2-6

Fri Jul 9 2021, Lucie Maňásková (Imanasko@redhat.com)

- Updated text for [JIRA:RHELPLAN-34199](#)(Security).

0.2-5

Wed Jun 23 2021, Lucie Maňásková (Imanasko@redhat.com)

- Added information about removal of AlternateTab in [BZ#1922488](#) (Desktop).

0.2-4

Fri May 21 2021, Lenka Špačková (lspackova@redhat.com)

- Updated information about OS conversion in [Overview](#).

0.2-3

Thu May 20 2021, Lenka Špačková (lspackova@redhat.com)

- Added a workaround to the known issue [BZ#1942330](#) (Dynamic programming languages, web and database servers).

0.2-2

Fri May 14 2021, Lucie Maňásková (Imanasko@redhat.com)

- Added a new feature [BZ#1944677](#) about .NET 5 support (Compilers and development tools).
- Added a new feature xfer:BZ-1959289[BZ#1959289] (RHEL System Roles).
- Updated information about the **xdp-tools** package in [BZ#1820670](#) (Networking) and in [BZ#1780124](#) (Kernel).

0.2-1

Mon Apr 19 2021, Lenka Špačková (lspackova@redhat.com)

- Added a known issue [BZ#1942330](#) (Dynamic programming languages, web and database servers).

0.2-0

Tue Apr 13 2021, Lenka Špačková (lspackova@redhat.com)

- Added a known issue (Installer and image creation).

0.1-9

Tue Apr 06 2021, Lenka Špačková (lspackova@redhat.com)

- Improved the list of supported architectures.

0.1-8

Wed Mar 31 2021, Lenka Špačková (lspackova@redhat.com)

- Updated information about OS conversions with the availability of the supported **Convert2RHEL** utility.

0.1-7

Mon Mar 29 2021, Lucie Maňásková (Imanasko@redhat.com)

- Updated the New features section (Kernel).

0.1-6

Thu Feb 25 2021, Lenka Špačková (lspackova@redhat.com)

- Fixed CentOS Linux name.

0-1-5

Tue Feb 23 2021, Lucie Maňásková (Imanasko@redhat.com)

- Add a known issues (Identity Management).
- Add a note about the **podman** utility rebase to the RHEL 8.3.1 section.

0-1-4

Thu Feb 18 2021, Jaroslav Klech (jklech@redhat.com)

- Adds a known issue (Kernel).
- Fix links for an enhancement (Kernel).

0-1-3

Tue Feb 16 2021, Lenka Špačková (lspackova@redhat.com)

- Release of the Red Hat Enterprise Linux 8.3.1 Release Notes.
- Update in-place upgrade section in Overview with the release of the [RHBA-2021:0569](#) advisory.

0-1-2

Fri Feb 12 2021, Lucie Maňásková (Imanasko@redhat.com)

- Added two known issues (Security, Installer).

0-1-1

Wed Feb 10 2021, Lucie Maňásková (Imanasko@redhat.com)

- Added a known issue (Virtualization).

0-1-0

Wed Feb 03 2021, Lenka Špačková (lspackova@redhat.com)

- Added a note about the consolidation of network configuration in the kernel command line under the **ip** parameter (Networking).
- Added **mercurial** to deprecated packages.

- Added a known issue related to Witherspoon hosts (Virtualization).

0-0-9

Fri Jan 29 2021, Lucie Maňásková (lmanasko@redhat.com)

- Added new bug fix description (Security).
- Added a note about deprecation of the **mailman** package (Software management).
- Updated the New features section (Security, Identity Management).
- Added Technology Preview note about the **systemd-resolved** service.
- Other minor updates.

0.0-8

Mon Dec 14 2020, Lucie Maňásková (lmanasko@redhat.com)

- Updated the Known issues section and the Bug fixes section.

0.0-7

Fri Nov 27 2020, Lucie Maňásková (lmanasko@redhat.com)

- Added a bug fix for issue with **fapolicyd** (Security).
- More updates to the Bug Fixes section.
- Added a note about deprecation of the Podman varlink-based REST API V1 (Containers).
- Updated the New features section.
- Added new Known issue about replicating blueprints from the **lorax-composer** back end to the new **osbuild-composer** back end (Image Builder).

0.0-6

Fri Nov 20 2020, Lucie Maňásková (lmanasko@redhat.com)

- Added an OpenSCAP bug fix description (Security).
- Updated the New features section (Software management).

0.0-5

Wed Nov 18 2020, Lenka Špačková (lspackova@redhat.com)

- Added information about conversion from Oracle Linux or CentOS Linux to RHEL (Overview).

0.0-4

Thu Nov 12 2020, Lenka Špačková (lspackova@redhat.com)

- Added information about **Node.js 14.15.0** released with the [RHEA-2020:5101](#) advisory.

0.0-3

Wed Nov 11 2020, Lucie Maňásková (lmanasko@redhat.com)

- Added description about Omni-Path Architecture (OPA) host software support to New features.

0.0-2

Mon Nov 09 2020, Lenka Špačková (lspackova@redhat.com)

- Added Intel Tiger Lake graphics as a Technology Preview (Graphics infrastructures).

0.0-1

Wed Nov 04 2020, Lucie Maňásková ([lmanasko@redhat.com](mailto:لماناسكو@redhat.com))

- Release of the Red Hat Enterprise Linux 8.3 Release Notes.

0.0-0

Tue Jul 28 2020, Lucie Maňásková ([lmanasko@redhat.com](mailto:لماناسكو@redhat.com))

- Release of the Red Hat Enterprise Linux 8.3 Beta Release Notes.