



Red Hat Enterprise Linux 8

8.8 Release Notes

Release Notes for Red Hat Enterprise Linux 8.8

Red Hat Enterprise Linux 8 8.8 Release Notes

Release Notes for Red Hat Enterprise Linux 8.8

Legal Notice

Copyright © 2024 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 8.8 and document known problems in this release, as well as notable bug fixes, Technology Previews, deprecated functionality, and other details. For information about installing Red Hat Enterprise Linux, see Installation.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	5
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	6
CHAPTER 1. OVERVIEW	7
1.1. MAJOR CHANGES IN RHEL 8.8	7
Installer and image creation	7
RHEL for Edge	7
Security	7
Dynamic programming languages, web and database servers	7
Compilers and development tools	8
Updated performance tools and debuggers	8
Updated performance monitoring tools	8
Updated compiler toolsets	8
Java implementations in RHEL 8	8
The web console	8
Containers	9
1.2. IN-PLACE UPGRADE AND OS CONVERSION	9
In-place upgrade from RHEL 7 to RHEL 8	9
In-place upgrade from RHEL 6 to RHEL 8	10
In-place upgrade from RHEL 8 to RHEL 9	10
Conversion from a different Linux distribution to RHEL	10
1.3. RED HAT CUSTOMER PORTAL LABS	10
1.4. ADDITIONAL RESOURCES	11
CHAPTER 2. ARCHITECTURES	12
CHAPTER 3. DISTRIBUTION OF CONTENT IN RHEL 8	13
3.1. INSTALLATION	13
3.2. REPOSITORIES	13
3.3. APPLICATION STREAMS	14
3.4. PACKAGE MANAGEMENT WITH YUM/DNF	14
CHAPTER 4. NEW FEATURES	15
4.1. INSTALLER AND IMAGE CREATION	15
4.2. RHEL FOR EDGE	15
4.3. SOFTWARE MANAGEMENT	16
4.4. SHELLS AND COMMAND-LINE TOOLS	17
4.5. INFRASTRUCTURE SERVICES	17
4.6. SECURITY	18
4.7. NETWORKING	22
4.8. KERNEL	24
4.9. HIGH AVAILABILITY AND CLUSTERS	27
4.10. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	27
4.11. COMPILERS AND DEVELOPMENT TOOLS	33
4.12. IDENTITY MANAGEMENT	36
4.13. DESKTOP	40
4.14. THE WEB CONSOLE	41
4.15. RED HAT ENTERPRISE LINUX SYSTEM ROLES	42
4.16. VIRTUALIZATION	47
4.17. SUPPORTABILITY	48
4.18. CONTAINERS	48

CHAPTER 5. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS	52
New kernel parameters	52
Updated kernel parameters	52
New sysctl parameters	55
CHAPTER 6. DEVICE DRIVERS	56
6.1. NEW DRIVERS	56
Network drivers	56
Graphics drivers and miscellaneous drivers	56
6.2. UPDATED DRIVERS	56
Network drivers	56
Storage drivers	57
CHAPTER 7. AVAILABLE BPF FEATURES	58
CHAPTER 8. BUG FIXES	72
8.1. INSTALLER AND IMAGE CREATION	72
8.2. SOFTWARE MANAGEMENT	73
8.3. SHELLS AND COMMAND-LINE TOOLS	73
8.4. INFRASTRUCTURE SERVICES	74
8.5. SECURITY	74
8.6. NETWORKING	80
8.7. KERNEL	81
8.8. FILE SYSTEMS AND STORAGE	81
8.9. HIGH AVAILABILITY AND CLUSTERS	81
8.10. COMPILERS AND DEVELOPMENT TOOLS	83
8.11. IDENTITY MANAGEMENT	83
8.12. GRAPHICS INFRASTRUCTURE	84
8.13. THE WEB CONSOLE	85
8.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES	85
8.15. VIRTUALIZATION	86
CHAPTER 9. TECHNOLOGY PREVIEWS	88
9.1. INFRASTRUCTURE SERVICES	88
9.2. NETWORKING	88
9.3. KERNEL	89
9.4. FILE SYSTEMS AND STORAGE	91
9.5. HIGH AVAILABILITY AND CLUSTERS	93
9.6. IDENTITY MANAGEMENT	94
9.7. DESKTOP	96
9.8. GRAPHICS INFRASTRUCTURES	97
9.9. VIRTUALIZATION	97
9.10. RHEL IN CLOUD ENVIRONMENTS	99
9.11. CONTAINERS	99
CHAPTER 10. DEPRECATED FUNCTIONALITY	101
10.1. INSTALLER AND IMAGE CREATION	101
10.2. SUBSCRIPTION MANAGEMENT	102
10.3. SOFTWARE MANAGEMENT	102
10.4. SHELLS AND COMMAND-LINE TOOLS	102
10.5. SECURITY	104
10.6. NETWORKING	105
10.7. KERNEL	107
10.8. BOOT LOADER	108

10.9. FILE SYSTEMS AND STORAGE	108
10.10. HIGH AVAILABILITY AND CLUSTERS	109
10.11. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	110
10.12. COMPILERS AND DEVELOPMENT TOOLS	110
10.13. IDENTITY MANAGEMENT	111
10.14. DESKTOP	114
10.15. GRAPHICS INFRASTRUCTURES	114
10.16. THE WEB CONSOLE	115
10.17. RED HAT ENTERPRISE LINUX SYSTEM ROLES	115
10.18. VIRTUALIZATION	116
10.19. CONTAINERS	117
10.20. DEPRECATED PACKAGES	118
10.21. DEPRECATED AND UNMAINTAINED DEVICES	155
CHAPTER 11. KNOWN ISSUES	159
11.1. INSTALLER AND IMAGE CREATION	159
11.2. SUBSCRIPTION MANAGEMENT	161
11.3. SOFTWARE MANAGEMENT	161
11.4. SHELLS AND COMMAND-LINE TOOLS	162
11.5. INFRASTRUCTURE SERVICES	163
11.6. SECURITY	163
11.7. NETWORKING	169
11.8. KERNEL	169
11.9. BOOT LOADER	175
11.10. FILE SYSTEMS AND STORAGE	176
11.11. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	177
11.12. IDENTITY MANAGEMENT	178
11.13. DESKTOP	181
11.14. GRAPHICS INFRASTRUCTURES	182
11.15. THE WEB CONSOLE	183
11.16. RED HAT ENTERPRISE LINUX SYSTEM ROLES	184
11.17. VIRTUALIZATION	184
11.18. RHEL IN CLOUD ENVIRONMENTS	188
11.19. SUPPORTABILITY	190
11.20. CONTAINERS	191
CHAPTER 12. INTERNATIONALIZATION	192
12.1. RED HAT ENTERPRISE LINUX 8 INTERNATIONAL LANGUAGES	192
12.2. NOTABLE CHANGES TO INTERNATIONALIZATION IN RHEL 8	192
APPENDIX A. LIST OF TICKETS BY COMPONENT	194
APPENDIX B. REVISION HISTORY	202

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your feedback on our documentation. Let us know how we can improve it.

Submitting feedback through Jira (account required)

1. Log in to the [Jira](#) website.
2. Click **Create** in the top navigation bar.
3. Enter a descriptive title in the **Summary** field.
4. Enter your suggestion for improvement in the **Description** field. Include links to the relevant parts of the documentation.
5. Click **Create** at the bottom of the dialogue.

CHAPTER 1. OVERVIEW

1.1. MAJOR CHANGES IN RHEL 8.8

Installer and image creation

Key highlights for image builder:

- Image builder on-prem now offers a new and improved way to create blueprints and images in the image builder web console.
- The RHEL for Edge Simplified Installer image type is now available in the image builder web console.

For more information, see [New features - Installer and image creation](#).

RHEL for Edge

RHEL for Edge introduces the following new feature in RHEL 8.8:

- Specifying a user in a blueprint for **simplified-installer** images is now supported.

For more information, see [New features - RHEL for Edge](#).

Security

Key security-related highlights:

- The **FIPS mode** settings in the kernel have been adjusted to conform to the Federal Information Processing Standard (FIPS) 140-3. This change introduces stricter settings to many cryptographic algorithms, functions, and cipher suites.
- The **Libreswan** IPsec implementation was rebased to version 4.9.
- With the **fapolicyd** software framework, you can now filter the RPM database.
- The **OpenSCAP** security compliance utility was rebased to version 1.3.7.
- **Rsyslog** TLS-encrypted logging now supports multiple CA files.
- The **systemd-socket-proxyd** service now runs in its own SELinux domain due to an update to the SELinux policy.

See [New features - Security](#) for more information.

Dynamic programming languages, web and database servers

Later versions of the following Application Streams are now available:

- **Python 3.11**
- **nginx 1.22**
- **PostgreSQL 15**

The following components have been upgraded:

- **Git** to version 2.39.1
- **Git LFS** to version 3.2.0

See [New features - Dynamic programming languages, web and database servers](#) for more information.

Compilers and development tools

Updated performance tools and debuggers

The following performance tools and debuggers have been updated in RHEL 8.8:

- **Valgrind 3.19**
- **SystemTap 4.8**
- **elfutils 0.188**

Updated performance monitoring tools

The following performance monitoring tools have been updated in RHEL 8.8:

- **PCP 5.3.7**
- **Grafana 7.5.15**

Updated compiler toolsets

The following compiler toolsets have been updated in RHEL 8.8:

- **GCC Toolset 12**
- **LLVM Toolset 15.0.7**
- **Rust Toolset 1.66**
- **Go Toolset 1.19.4**

See [New features - Compilers and development tools](#) for more information.

Java implementations in RHEL 8

The RHEL 8 AppStream repository includes:

- The **java-17-openjdk** packages, which provide the OpenJDK 17 Java Runtime Environment and the OpenJDK 17 Java Software Development Kit.
- The **java-11-openjdk** packages, which provide the OpenJDK 11 Java Runtime Environment and the OpenJDK 11 Java Software Development Kit.
- The **java-1.8.0-openjdk** packages, which provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Software Development Kit.

The Red Hat build of OpenJDK packages share a single set of binaries between its portable Linux releases, RHEL 8.8 and later releases. Because of this update, there is a change in the process of rebuilding the OpenJDK packages on RHEL from the source RPM. For more information about the new rebuilding process, see the **README.md** file which is available in the SRPM package of the Red Hat build of OpenJDK and is also installed by the **java-*-openjdk-headless** packages under the **/usr/share/doc** tree.

For more information, see [OpenJDK documentation](#).

The web console

The RHEL web console now performs additional steps for binding LUKS-encrypted root volumes to **NBDE** deployments.

You can also apply the following **cryptographic subpolicies** through the graphical interface now: **DEFAULT:SHA1**, **LEGACY:AD-SUPPORT**, and **FIPS:OSPP**.

See [New features - The web console](#) for more information.

Containers

Notable changes include:

- The **podman** RHEL System Role is now available.
- Clients for sigstore signatures with Fulcio and Rekor are now available.
- Skopeo now supports generating sigstore key pairs.
- Podman now supports events for auditing.
- The Container Tools packages have been updated.
- The Aardvark and Netavark networks stack now supports custom DNS server selection.
- Toolbox is now available.
- Podman Quadlet is now available as a Technology Preview.
- The **container-tools:3.0** module stream has been deprecated.
- The CNI network stack has been deprecated.

See [New features - Containers](#) for more information.

1.2. IN-PLACE UPGRADE AND OS CONVERSION

In-place upgrade from RHEL 7 to RHEL 8

The possible in-place upgrade paths currently are:

- From RHEL 7.9 to RHEL 8.6 and RHEL 8.8 on the 64-bit Intel, IBM POWER 8 (little endian), and IBM Z architectures
- From RHEL 7.6 to RHEL 8.4 on architectures that require kernel version 4.14: IBM POWER 9 (little endian) and IBM Z (Structure A). This is the final in-place upgrade path for these architectures.
- From RHEL 7.9 to RHEL 8.6 and RHEL 8.8 on systems with SAP HANA on the 64-bit Intel architecture.

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](#) .

If you are upgrading to RHEL 8.8 with SAP HANA, ensure that the system is certified for SAP prior to the upgrade. For instructions on performing an in-place upgrade on systems with SAP environments, see [How to in-place upgrade SAP environments from RHEL 7 to RHEL 8](#) .



NOTE

For the successful in-place upgrade of RHEL 7.6 for IBM POWER 9 (little endian) and IBM Z (structure A) architectures, you must manually download the specific Leapp data. For more information, see the [Leapp data snapshots for an in-place upgrade](#) Knowledgebase article.

Notable enhancements include:

- The RHEL in-place upgrade path strategy has changed. For more information, see [Supported in-place upgrade paths for Red Hat Enterprise Linux](#).
- The latest release of the **leapp-upgrade-el7toel8** package now contains all required data files. Customers no longer need to manually download these data files.
- In-place upgrades using an ISO image that contains the target version are now possible.
- RPM signatures are now automatically checked during the in-place upgrade. To disable the automatic check, use the **--nogpgcheck** option when performing the upgrade.
- Systems that are subscribed to RHSM are now automatically registered with Red Hat Insights. To disable the automatic registration, set the **LEAPP_NO_INSIGHTS_REGISTER** environment variable to **1**.
- Red Hat now collects upgrade-related data, such as the upgrade start and end times and whether the upgrade was successful, for utility usage analysis. To disable data collection, set the **LEAPP_NO_RHSM_FACTS** environment variable to **1**.

In-place upgrade from RHEL 6 to RHEL 8

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

In-place upgrade from RHEL 8 to RHEL 9

Instructions on how to perform an in-place upgrade from RHEL 8 to RHEL 9 using the Leapp utility are provided by the document [Upgrading from RHEL 8 to RHEL 9](#). Major differences between RHEL 8 and RHEL 9 are documented in [Considerations in adopting RHEL 9](#).

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 perform an unsupported conversion from a RHEL-derived Linux distribution to RHEL](#).

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

1.3. 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](#)
- [Red Hat Memory Analyzer](#)
- [Kernel Oops Analyzer](#)
- [Red Hat Product Errata Advisory Checker](#)

1.4. 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](#).
- 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** including removed functionality, 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.

CHAPTER 2. ARCHITECTURES

Red Hat Enterprise Linux 8.8 is distributed with the kernel version 4.18.0-477.10, 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 Installation ISO image is in multiple GB size, and as a result, it might not fit on optical media formats. A USB key or USB hard drive is recommended when using the Installation 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. NEW FEATURES

This part describes new features and major enhancements introduced in Red Hat Enterprise Linux 8.8.

4.1. INSTALLER AND IMAGE CREATION

A new and improved way to create blueprints and images in the image builder web console

With this enhancement, you have access to a unified version of the image builder tool and a significant improvement in your user experience.

Notable enhancements in the image builder dashboard GUI include:

- You can now customize your blueprints with all the customizations previously supported only in the CLI, such as kernel, file system, firewall, locale, and other customizations.
- You can import blueprints by either uploading or dragging the blueprint in the **.JSON** or **.TOML** format and create images from the imported blueprint.
- You can also export or save your blueprints in the **.JSON** or **.TOML** format.
- Access to a blueprint list that you can sort, filter, and is case-sensitive.
- With the image builder dashboard, you can now access your blueprints, images, and sources by navigating through the following tabs:
 - Blueprint - Under the Blueprint tab, you can now import, export, or delete your blueprints.
 - Images - Under the Images tab, you can:
 - Download images.
 - Download image logs.
 - Delete images.
 - Sources - Under the Sources tab, you can:
 - Download images.
 - Download image logs.
 - Create sources for images.
 - Delete images.

Jira:RHELPLAN-139448

Support for 64-bit ARM for **.vhd** images built with image builder

Previously, Microsoft Azure **.vhd** images created with the image builder tool were not supported on 64-bit ARM architectures. This update adds support for 64-bit ARM Microsoft Azure **.vhd** images and now you can build your **.vhd** images using image builder and upload them to the Microsoft Azure cloud.

Jira:RHELPLAN-139424

4.2. RHEL FOR EDGE

Ability to specify user in a blueprint for **simplified-installer** images

Previously, when creating a blueprint for a **simplified-installer** image, you could not specify a user in the blueprint customization, because the customization was not used and was discarded. With this update, when you create an image from the blueprint, this blueprint creates a user under the **/usr/lib/passwd** directory and a password under the **/usr/etc/shadow** directory during installation time. You can log in to the device with the username and the password you created for the blueprint. Note that after you access the system, you need to create users, for example, using the **useradd** command.

Jira:RHELPLAN-149091

Red Hat build of MicroShift enablement for RHEL for Edge images

With this enhancement, you can enable Red Hat build of MicroShift services in a RHEL for Edge system. By using the **[[customizations.firewall.zones]]** blueprint customization, you can add support for **firewalld** sources in the blueprint customization. For that, specify a name for the zone and a list of sources in that specific zone. Sources can be of the form **source[/mask][MAC]ipset:ipset**.

The following is a blueprint example on how to configure and customize support for Red Hat build of MicroShift services in a RHEL for Edge system.

```
[[packages]]
name = "microshift"
version = "*"
[customizations.services]
enabled = ["microshift"]
[[customizations.firewall.zones]]
name = "trusted"
sources = ["10.42.0.0/16", "169.254.169.1"]
```


The Red Hat build of MicroShift installation requirements, such as firewall policies, MicroShift RPM, **systemd** service, enable you to create a deployment ready for production to achieve workload portability to a minimum field deployed edge device and by default LVM device mapper enablement.

Jira:RHELPLAN-136489

4.3. SOFTWARE MANAGEMENT

New **yum offline-upgrade** command for offline updates on RHEL

With this enhancement, you can apply offline updates to RHEL by using the new **yum offline-upgrade** command from the YUM **system-upgrade** plug-in.



IMPORTANT

The **yum system-upgrade** command included in the **system-upgrade** plug-in is not supported on RHEL.

[Bugzilla:2054235](#)

Applying advisory security filters to **yum offline-upgrade** is now supported

With this enhancement, the new functionality for advisories filtering has been added. As a result, you can now download packages and their dependencies only from the specified advisory by using the **yum offline-upgrade** command with advisory security filters (**--advisory**, **--security**, **--bugfix**, and other filters).

[Bugzilla:2139324](#)

The `unload_plugins` function is now available for the YUM API

With this enhancement, a new **unload_plugins** function has been added to the YUM API to allow plugins unloading.



IMPORTANT

Note that you must first run the **init_plugins** function, and then run the **unload_plugins** function.

[Bugzilla:2047251](#)

New `--nocompression` option for `rpm2archive`

With this enhancement, the **--nocompression** option has been added to the **rpm2archive** utility. You can use this option to avoid compression when directly unpacking an RPM package.

[Bugzilla:2129345](#)

4.4. SHELLS AND COMMAND-LINE TOOLS

ReaR is now fully supported also on the 64-bit IBM Z architecture

Basic Relax and Recover (ReaR) functionality, previously available on the 64-bit IBM Z architecture as a Technology Preview, is fully supported with the **rear** package version 2.6-9.el8 or later. You can create a ReaR rescue image on the IBM Z architecture in the z/VM environment only. Backing up and recovering logical partitions (LPARs) is not supported at the moment. ReaR supports saving and restoring disk layout only on Extended Count Key Data (ECKD) direct access storage devices (DASDs). Fixed Block Access (FBA) DASDs and SCSI disks attached through Fibre Channel Protocol (FCP) are not supported for this purpose. The only output method currently available is Initial Program Load (IPL), which produces a kernel and an initial ramdisk (initrd) compatible with the **zipl** bootloader.

For more information see [Using a ReaR rescue image on the 64-bit IBM Z architecture](#) .

[Bugzilla:2130206](#), [Bugzilla:1868421](#)

4.5. INFRASTRUCTURE SERVICES

New `synce4l` package for frequency synchronization is now available

SyncE (Synchronous Ethernet) is a hardware feature that enables PTP clocks to achieve precise synchronization of frequency at the physical layer. SyncE is supported in certain network interface cards (NICs) and network switches.

With this enhancement, the new **synce4l** package is now available, which provides support for SyncE. As a result, Telco Radio Access Network (RAN) applications can now achieve more efficient communication due to more accurate time synchronization.

[Bugzilla:2019751](#)

`powertop` rebased to version 2.15

The **powertop** package for improving the energy efficiency has been updated to version 2.15. Notable changes and enhancements include:

- Several Valgrind errors and possible buffer overrun have been fixed to improve the **powertop** tool stability.
- Improved compatibility with Ryzen processors and Kaby Lake platforms.
- Enabled Lake Field, Alder Lake N, and Raptor Lake platforms support.
- Enabled Ice Lake NNPI and Meteor Lake mobile and desktop support.

Bugzilla:2040070

tuned rebased to version 2.20.0

The TuneD utility for optimizing the performance of applications and workloads has been updated to version 2.20.0. Notable changes and enhancements over version 2.19.0 include:

- An extension of API enables you to move devices between plug-in instances at runtime.
- The **plugin_cpu** module, which provides fine-tuning of CPU-related performance settings, introduces the following enhancements:
 - The **pm_qos_resume_latency_us** feature enables you to limit the maximum time allowed for each CPU to transition from an idle state to an active state.
 - TuneD adds support for the **intel_pstate** scaling driver, which provides scaling algorithms to tune the systems' power management based on different usage scenarios.
- The socket API to control TuneD through a Unix domain socket is now available as a Technology Preview. See [Socket API for TuneD available as a Technology Preview](#) for more information.

[Bugzilla:2133814](#), [Bugzilla:2113925](#), [Bugzilla:2118786](#), [Bugzilla:2095829](#), [Bugzilla:2113900](#)

4.6. SECURITY

FIPS mode now has more secure settings that target FIPS 140-3

The FIPS mode settings in the kernel have been adjusted to conform to the Federal Information Processing Standard (FIPS) 140-3. This change introduces stricter settings to many cryptographic algorithms, functions, and cipher suites. Most notably:

- The Triple Data Encryption Standard (3DES), Elliptic-curve Diffie-Hellman (ECDH), and Finite-Field Diffie-Hellman (FFDH) algorithms are now disabled. This change affects Bluetooth, DH-related operations in the kernel keyring, and Intel QuickAssist Technology (QAT) cryptographic accelerators.
- The hash-based message authentication code (HMAC) key now cannot be shorter than 112 bits. The minimum key length is set to 2048 bits for Rivest-Shamir-Adleman (RSA) algorithms.
- Drivers that used the **xts_check_key()** function have been updated to use the **xts_verify_key()** function instead.
- The following Deterministic Random Bit Generator (DRBG) hash functions are now disabled: SHA-224, SHA-384, SHA512-224, SHA512-256, SHA3-224, and SHA3-384.



NOTE

Even though the RHEL 8.6 (and newer) kernel in FIPS mode is designed to be compliant with FIPS 140-3, it is not yet certified by the National Institute of Standards and Technology (NIST) Cryptographic Module Validation Program (CMVP). The latest certified kernel module is the updated RHEL 8.5 kernel after the RHSA-2021:4356 advisory update. That certification applies to the FIPS 140-2 standard. You cannot choose whether a cryptographic module conforms to FIPS 140-2 or 140-3. For more information, see the [Compliance Activities and Government Standards: FIPS 140-2 and FIPS 140-3](#) Knowledgebase article.

Bugzilla:2107595, Bugzilla:2158893, Bugzilla:2175234, Bugzilla:2166715, Bugzilla:2129392,
[Bugzilla:2152133](#)

Libreswan rebased to 4.9

The **libreswan** packages have been upgraded to version 4.9. Notable changes over the previous version include:

- Added support for **{left,right}pubkey=** to the **addconn** and **whack** utilities
- Added key derivation function (KDF) self-tests
- Updated list of allowed system calls for the **seccomp** filter
- Show host's authentication key (**showhostkey**):
 - Added support for Elliptic Curve Digital Signature Algorithm (ECDSA) pubkeys
 - Added the **--pem** option to print Privacy-Enhanced Mail (PEM)-encoded public key
- The Internet Key Exchange Protocol Version 2 (IKEv2):
 - Extensible Authentication Protocol – Transport Layer Security (EAP-TLS) support
 - EAP-only authentication support
 - Labeled IPsec improvements
- The **pluto** Internet Key Exchange (IKE) daemon:
 - Support for **maxbytes** and **maxpacket** counters
 - Changed default value of **replay-window** from 32 to 128
 - Changed the default value of **esn=** to **either** and preferred value to **yes**
 - Disabled **esn** when **replay-window=** is set to **0**
 - Dropped obsolete debug options such as **crypto-low**

Bugzilla:2128672

SELinux now confines udftools

With this update of the **selinux-policy** packages, SELinux confines the **udftools** service.

Bugzilla:1972230

New SELinux policy for `systemd-socket-proxyd`

Because the **systemd-socket-proxyd** service requires particular resources usage, a new policy with the required rules was added to the **selinux-policy** packages. As a result, the service runs in its SELinux domain.

[Bugzilla:2088441](#)

OpenSCAP rebased to 1.3.7

The OpenSCAP packages have been rebased to upstream version 1.3.7. This version provides various bug fixes and enhancements, most notably:

- Fixed error when processing OVAL filters (rhbz#2126882)
- OpenSCAP no longer emits invalid empty **xmlfilecontent** items if XPath does not match (rhbz#2139060)
- Prevented **Failed to check available memory** errors (rhbz#2111040)

[Bugzilla:2159290](#)

scap-security-guide rules for Rsyslog log files are compatible with RainerScript

Rules in **scap-security-guide** for checking and remediating ownership, group ownership, and permissions of Rsyslog log files are now also compatible with log files defined by using the RainerScript syntax. Modern systems already use the RainerScript syntax in Rsyslog configuration files and the respective rules were not able to recognize this syntax. As a result, **scap-security-guide** rules can now check and remediate ownership, group ownership, and permissions of Rsyslog log files in both available syntaxes.

[Bugzilla:2072444](#)

STIG security profile updated to version V1R9

The **DISA STIG for Red Hat Enterprise Linux 8** profile in the SCAP Security Guide has been updated to align with the latest version **V1R9**. This release also includes changes published in **V1R8**.

Use only the current version of this profile because previous versions are no longer valid.

The following STIG IDs have been updated:

- V1R9
 - RHEL-08-010359 - Selected rule **aide_build_database**
 - RHEL-08-010510 - Removed rule **sshd_disable_compression**
 - RHEL-08-020040 - New rule to configure tmux keybinding
 - RHEL-08-020041 - New rule to configure starting **tmux** instead of **exec tmux**
- V1R8
 - Multiple STIG IDs - The **sshd** and **sysctl** rules can identify and remove duplicate or conflicting configurations.
 - RHEL-08-010200 - SSHD ClientAliveCountMax is configured with value **1**.

- RHEL-08-020352 - Check and remediations now ignore **.bash_history**.
- RHEL-08-040137 - Check updated to examine both **/etc/fapolicyd/fapolicyd.rules** and **/etc/fapolicyd/complied.rules**.



WARNING

Automatic remediation might make the system non-functional. Run the remediation in a test environment first.

[Bugzilla:2152658](#)

RHEL 8 STIG profiles are better aligned with the benchmark

Four existing rules that satisfy RHEL 8 STIG requirements were part of the data stream but were previously not included in the STIG profiles (**stig** and **stig_gui**). With this update, the following rules are now included in the profiles:

- **accounts_passwords_pam_faillock_dir**
- **accounts_passwords_pam_faillock_silent**
- **account_password_selinux_faillock_dir**
- **fapolicy_default_deny**

As a result, the RHEL 8 STIG profiles have a higher coverage.

[Bugzilla:2156192](#)

SCAP Security Guide rebased to 0.1.66

The SCAP Security Guide (SSG) packages have been rebased to upstream version 0.1.66. This version provides various enhancements and bug fixes, most notably:

- Updated RHEL 8 STIG profiles
- Deprecated rule **account_passwords_pam_faillock_audit** in favor of **accounts_passwords_pam_faillock_audit**

[Bugzilla:2158404](#)

OpenSSL driver can now use certificate chains in Rsyslog

The **NetstreamDriverCaExtraFiles** directive allows configuring multiple additional certificate authority (CA) files. With this update, you can specify multiple CA files and the OpenSSL library can validate them, which is necessary for SSL certificate chains. As a result, you can use certificate chains in Rsyslog with the OpenSSL driver.

[Bugzilla:2124934](#)

opencryptoki rebased to 3.19.0

The **opencryptoki** package has been rebased to version 3.19.0, which provides many enhancements and bug fixes. Most notably, **opencryptoki** now supports the following features:

- IBM-specific Dilithium keys
- Dual-function cryptographic functions
- Cancelling active session-based operations by using the new **C_SessionCancel** function, as described in the PKCS #11 Cryptographic Token Interface Base Specification v3.0
- Schnorr signatures through the **CKM_IBM_ECDSA_OTHER** mechanism
- Bitcoin key derivation through the **CKM_IBM_BTC_DERIVE** mechanism
- EP11 tokens in IBM z16 systems

Bugzilla:2110315

New SCAP rule for idle session termination

New SCAP rule **logind_session_timeout** has been added to the **scap-security-guide** package in ANSSI-BP-028 profiles for Enhanced and High levels. This rule uses a new feature of the **systemd** service manager and terminates idle user sessions after a certain time. This rule provides automatic configuration of a robust idle session termination mechanism which is required by multiple security policies. As a result, OpenSCAP can automatically check the security requirement related to terminating idle user sessions and, if necessary, remediate it.

[Bugzilla:2122322](#)

fapolicyd now provides filtering of the RPM database

With the new configuration file **/etc/fapolicyd/rpm-filter.conf**, you can customize the list of RPM-database files that the **fapolicyd** software framework stores in the trust database. This way, you can block certain applications installed by RPM or allow an application denied by the default configuration filter.

[Bugzilla:2165645](#)

4.7. NETWORKING

The default MPTCP subflow limit is 2

A subflow is a single TCP connection that is part of a Multipath TCP (MPTCP) connection. A subflow limit in MPTCP refers to the maximum number of additional connections that can be created between two MPTCP endpoints. You can use the limit to restrict the number of additional parallel subflows that can be created between the endpoints, to avoid overloading the network and the endpoints. For example the value of 0 allows only the initial subflow.

With this enhancement, the default MPTCP subflow limit has been increased from 0 to 2. This enables you by default to create multiple additional subflows. If you need a different value, you can create a Systemd oneshot unit. The unit should execute the **ip mptcp limits set subflows <YOUR_VALUE>** command after your network (**network.target**) is operational during every boot process.

Bugzilla:2127136

The kernel now logs the listening address in SYN flood messages

This enhancement adds the listening IP address to SYN flood messages:


Possible SYN flooding on port <ip_address>:<port>.

As a result, if many processes are bound to the same port on different IP addresses, administrators can now clearly identify the affected socket.

Bugzilla:2143849

The **nm-initrd-generator** profiles now have lower priority than autoconnect profiles

The **nm-initrd-generator** early boot NetworkManager configuration generator utility generates and configures connection profiles by using the NetworkManager instance running in the boot loader's initialized **initrd** RAM disk. The **nm-initrd-generator** utility generated profiles now have a lower autoconnect priority than the default connection autoconnect priority. This enables generated network profiles in **initrd** to coexist with user configuration in default root account.



NOTE

After switching from **initrd** root account to default root, the same profile stays activated and no new autoconnect happens.

Bugzilla:2089707

nispor rebased to version 1.2.10

The **nispor** packages have been upgraded to upstream version 1.2.10, which provides a number of enhancements and bug fixes over the previous version:

- Added support for **NetStateFilter** to use the kernel filter on network routes and interfaces.
- Single Root Input and Output Virtualization (SR-IOV) interfaces can query SR-IOV Virtual Function (SR-IOV VF) information per (VF).
- Newly supported bonding options: **lACP_active**, **arp_missed_max**, and **ns_ip6_target**.

Bugzilla:2153166

NetworkManager rebased to version 1.40.16

The **NetworkManager** packages have been upgraded to upstream version 1.40.16, which provides a number of bug fixes over the previous version:

- The **nm-cloud-setup** utility preserves externally added addresses.
- A race condition was fixed that prevented the automatic activation of MACsec connections at boot.
- NetworkManager now correctly calculates expiration times for items configured from IPv6 neighbor discovery messages.
- NetworkManager now automatically updates the **/etc/resolv.conf** file when the configuration changes.
- NetworkManager no longer sets non-existent interfaces as primary when activating a bond.
- Setting a primary interface in a bond now always works, even if the interface does not exist when you active the bond.

- The **NetworkManager --print-config** command no longer prints duplicate entries.
- The **ifcfg-rh** plug-in can now read InfiniBand P-Key connection profiles without an explicit interface name.
- The **nmcli** utility can now remove a bond port connection profile from a bond.
- A race condition was fixed that could occur during the activation of **veth** profiles if the peer already existed.
- NetworkManager now rejects DHCPv6 leases if all addresses fail IPv6 duplicate address detection (DAD).
- NetworkManager now waits until interfaces are connected before trying to resolve the system hostname on these interfaces from DNS.
- Profiles created by the **nm-initrd-generator** utility now have a lower-than-default priority.

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

[Bugzilla:2134907](#)

4.8. KERNEL

Kernel version in RHEL 8.8

Red Hat Enterprise Linux 8.8 is distributed with the kernel version 4.18.0-477.10.

[Bugzilla:2177769](#)

Secure Execution guest dump encryption with customer keys

This new feature allows Secure Execution guests to use hypervisor-initiated dumps to collect kernel crash information from KVM when the **kdump** utility does not work. Note that hypervisor-initiated dumps for Secure Execution are designed for the IBM Z Series z16 and LinuxONE Emperor 4 hardware.

[Bugzilla:2043833](#)

The **sfc** driver has split into **sfc** and **sfc_siena**

Following the changes in the upstream driver, the **sfc** NIC driver is now split into 2 different drivers: **sfc** and **sfc_siena**. **sfc_siena** supports the deprecated Siena family devices.

Note that custom configurations of the kernel module parameters and **udev** rules applied to **sfc** do not affect **sfc_siena** as they are now independent drivers. To customize both drivers, replicate the configuration options for **sfc_siena**.

[Bugzilla:2136107](#)

The **stmmac** driver is now fully supported

Red Hat now fully supports the **stmmac** driver for Intel® Elkhart Lake systems on a chip (SoCs).

[Bugzilla:1905243](#)

The **rtla** meta-tool adds the **osnoise** and **timerlat** tracers for improved tracer capabilities

The Real-Time Linux Analysis (**rtla**) is a meta-tool that includes a set of commands that analyze the

real-time properties of Linux. **rtla** leverages kernel tracing capabilities to provide accurate information about the properties and root causes of unexpected system results. **rtla** currently adds support for **osnoise** and **timerlat** tracer commands. The **osnoise** tracer reports a kernel thread per CPU. The **timerlat** tracer periodically prints the timer latency at the timer IRQ handler and the thread handler.

Note that to use the **timerlat** feature of **rtla**, you must disable admission control by using the **sysctl -w kernel.sched_rt_runtime_us=-1** script.

[Bugzilla:2075203](#)

The output format for **cgroups** and **irqs** has been improved to provide better readability

With this enhancement, the **tuna show_threads** command output for the **cgroup** utility is now structured based on the terminal size. You can also configure additional spacing to the **cgroups** output by adding the new **-z** or **--spaced** option to the **show_threads** command. As a result, you can now view the **cgroups** output in an improved readable format that is adaptable to your terminal size.

[Bugzilla:2121518](#)

The **rteval** command output now includes the program loads and measurement threads information

The **rteval** command now displays a report summary with the number of program loads, measurement threads, and the corresponding CPU that ran these threads. This information helps to evaluate the performance of a real-time kernel under load on specific hardware platforms.

The **rteval** report is written to an XML file along with the boot log for the system and saved to the **rteval-<date>-N.tar.bz2** compressed file. The **date** specifies the report generation date and **N** is the counter for the Nth run.

To generate an **rteval** report, enter the following command:

```
# rteval --summarize rteval-<date>-N.tar.bz2
```

[Bugzilla:2082260](#)

The **-W** and **--bucket-width** options have been added to the **oslat** program to measure latency

With this enhancement, you can specify a latency range for a single bucket at nanosecond accuracy. Widths that are not multiples of 1000 nanoseconds indicate nanosecond precision. By using the new options, **-W** or **--bucket-width**, you can modify the latency interval between buckets to measure latency within sub-microseconds delay time.

For example to set a latency bucket width of 100 nanoseconds for 32 buckets over a duration of 10 seconds to run on CPU range of 1-4 and omit zero bucket size, run the following command:

```
# oslat -b 32 -D 10s -W 100 -z -c 1-4
```

Note that before using the option, you must determine what level of precision is significant in relation to the error measurement.

[Bugzilla:2122374](#)

The Ethernet Port Configuration Tool (EPCT) utility support enabled in **E810** with Intel ice driver

With this enhancement, the **devlink port split** command now supports the Intel ice driver. The Ethernet Port Configuration Tool (EPCT) is a command line utility that allows you to change the link type of a device. The **devlink** utility, which displays device information and resources of devices, is dependent on EPCT. As a result of this enhancement, the ice driver implements support for EPCT, which enables you to list and view the configurable devices using Intel ice drivers.

Bugzilla:2009705

The Intel ice driver rebased to version 6.0.0

The Intel **ice** driver has been upgraded to upstream version 6.0.0, which provides a number of enhancements and bug fixes over previous versions. The notable enhancements include:

- Point-to-Point Protocol over Ethernet (**PPPoE**) protocol hardware offload
- Inter-Integrated Circuit (**I2C**) protocol write command
- VLAN Tag Protocol Identifier (**TPID**) filters in the Ethernet switch device driver model (**switchdev**)
- Double VLAN tagging in **switchdev**

Bugzilla:2103946

Hosting Secure Boot certificates for IBM zSystems

Starting with IBM z16 A02/AGZ and LinuxONE Rockhopper 4 LA2/AGL, you can manage certificates used to validate Linux kernels when starting the system with Secure Boot enabled on the Hardware Management Console (HMC). Notably:

- You can load certificates in a system certificate store using the HMC in DPM and classic mode from an FTP server that can be accessed by the HMC. It is also possible to load certificates from a USB device attached to the HMC.
- You can associate certificates stored in the certificate store with an LPAR partition. Multiple certificates can be associated with a partition and a certificate can be associated with multiple partitions.
- You can de-associate certificates in the certificate store from a partition by using HMC interfaces.
- You can remove certificates from the certificate store.
- You can associate up to 20 certificates with a partition.

The built-in firmware certificates are still available. In particular, as soon as you use the user-managed certificate store, the built-in certificates will no longer be available.

Certificate files loaded into the certificate store must meet the following requirements:

- They have the **PEM-** or **DER-encoded X.509v3** format and one of the following filename extensions: **.pem**, **.cer**, **.crt**, or **.der**.
- They are not expired.
- The key usage attribute must be *Digital Signature*.
- The extended key usage attribute must contain *Code Signing*.

A firmware interface allows a Linux kernel running in a logical partition to load the certificates associated with this partition. Linux on IBM Z stores these certificates in the **.platform** keyring, allowing the Linux kernel to verify **kexec** kernels and third party kernel modules to be verified using certificates associated with that partition.

It is the responsibility of the operator to only upload verified certificates and to remove certificates that have been revoked.



NOTE

The **Red Hat Secureboot 302** certificate that you need to load into the HMC is available at [Product Signing Keys](#).

Bugzilla:2183445

zipl support for Secure Boot IPL and dump on 64-bit IBM Z

With this update, the **zipl** utility supports List-Directed IPL and List-Directed dump from Extended Count Key Data (ECKD) Direct Access Storage Devices (DASD) on the 64-bit IBM Z architecture. As a result, Secure Boot for RHEL on IBM Z also works with the ECKD type of DASDs.

Bugzilla:2043852

4.9. HIGH AVAILABILITY AND CLUSTERS

New **enable-authfile** Booth configuration option

When you create a Booth configuration to use the Booth ticket manager in a cluster configuration, the **pcs booth setup** command now enables the new **enable-authfile** Booth configuration option by default. You can enable this option on an existing cluster with the **pcs booth enable-authfile** command. Additionally, the **pcs status** and **pcs booth status** commands now display warnings when they detect a possible **enable-authfile** misconfiguration.

[Bugzilla:2132582](#)

pcs can now run the **validate-all** action of resource and stonith agents

When creating or updating a resource or a STONITH device, you can now specify the **--agent-validation** option. With this option, **pcs** uses an agent's **validate-all** action, when it is available, in addition to the validation done by **pcs** based on the agent's metadata.

[Bugzilla:1816852](#), [Bugzilla:2159455](#)

4.10. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

Python 3.11 available in RHEL 8

RHEL 8.8 introduces Python 3.11, provided by the new package **python3.11** and a suite of packages built for it, as well as the **ubi8/python-311** container image.

Notable enhancements compared to the previously released Python 3.9 include:

- Significantly improved performance.

- Structural Pattern Matching using the new **match** keyword (similar to **switch** in other languages).
- Improved error messages, for example, indicating unclosed parentheses or brackets.
- Exact line numbers for debugging and other use cases.
- Support for defining context managers across multiple lines by enclosing the definitions in parentheses.
- Various new features related to type hints and the **typing** module, such as the new **X | Y** type union operator, variadic generics, and the new **Self** type.
- Precise error locations in tracebacks pointing to the expression that caused the error.
- A new **tomllib** standard library module which supports parsing TOML.
- An ability to raise and handle multiple unrelated exceptions simultaneously using Exception Groups and the new **except*** syntax.

Python 3.11 and packages built for it can be installed in parallel with Python 3.9, Python 3.8, and Python 3.6 on the same system.

Note that, unlike the previous versions, Python 3.11 is distributed as standard RPM packages instead of a module.

To install packages from the **python3.11** stack, use, for example:

```
# yum install python3.11
# yum install python3.11-pip
```

To run the interpreter, use, for example:

```
$ python3.11
$ python3.11 -m pip --help
```

See [Installing and using Python](#) for more information.

Note that Red Hat will continue to provide support for Python 3.6 until the end of life of RHEL 8. Similarly to Python 3.9, Python 3.11 will have a shorter life cycle; see [Red Hat Enterprise Linux Application Streams Life Cycle](#).

[Bugzilla:2137139](#)

nodejs:18 rebased to version 18.14 with npm rebased to version 9

Node.js 18.14, released in [RHSA-2023:1583](#), includes a SemVer major upgrade of **npm** from version 8 to version 9. This update was necessary due to maintenance reasons and may require you to adjust your **npm** configuration.

Notably, auth-related settings that are not scoped to a specific registry are no longer supported. This change was made for security reasons. If you used unscoped authentication configurations, the supplied token was sent to every registry listed in the **.npmrc** file.

If you use unscoped authentication tokens, generate and supply registry-scoped tokens in your **.npmrc** file.

If you have configuration lines using `_auth`, such as `//registry.npmjs.org/:_auth` in your `.npmrc` files, replace them with `//registry.npmjs.org/:_authToken=${NPM_TOKEN}` and supply the scoped token that you generated.

For a complete list of changes, see the [upstream changelog](#).

[Bugzilla:2178087](#)

git rebased to version 2.39.1

The **Git** version control system has been updated to version 2.39.1, which provides bug fixes, enhancements, and performance improvements over the previously released version 2.31.

Notable enhancements include:

- The **git log** command now supports a format placeholder for the **git describe** output: **git log --format=%(describe)**
- The **git commit** command now supports the **--fixup<commit>** option which enables you to fix the content of the commit without changing the log message. With this update, you can also use:
 - The **--fixup=amend:<commit>** option to change both the message and the content.
 - The **--fixup=rewrite:<commit>** option to update only the commit message.
- You can use the new **--reject-shallow** option with the **git clone** command to disable cloning from a shallow repository.
- The **git branch** command now supports the **--recurse-submodules** option.
- You can now use the **git merge-tree** command to:
 - Test if two branches can merge.
 - Compute a tree that would result in the merge commit if the branches were merged.
- You can use the new **safe.bareRepository** configuration variable to filter out bare repositories.

[Bugzilla:2139378](#)

git-lfs rebased to version 3.2.0

The **Git Large File Storage (LFS)** extension has been updated to version 3.2.0, which provides bug fixes, enhancements, and performance improvements over the previously released version 2.13.

Notable changes include:

- **Git LFS** introduces a pure SSH-based transport protocol.
- **Git LFS** now provides a merge driver.
- The **git lfs fsck** utility now additionally checks that pointers are canonical and that expected LFS files have the correct format.
- Support for the NT LAN Manager (NTLM) authentication protocol has been removed. Use Kerberos or Basic authentication instead.

[Bugzilla:2139382](#)

A new module stream: **nginx:1.22**

The **nginx 1.22** web and proxy server is now available as the **nginx:1.22** module stream. This update provides a number of bug fixes, security fixes, new features, and enhancements over the previously released version 1.20.

New features:

- **nginx** now supports:
 - OpenSSL 3.0 and the **SSL_sendfile()** function when using OpenSSL 3.0.
 - The PCRE2 library.
 - POP3 and IMAP pipelining in the **mail** proxy module.
- **nginx** now passes the **Auth-SSL-Protocol** and **Auth-SSL-Cipher** header lines to the mail proxy authentication server.

Enhanced directives:

- Multiple new directives are now available, such as **ssl_conf_command** and **ssl_reject_handshake**.
- The **proxy_cookie_flags** directive now supports variables.
- **nginx** now supports variables in the following directives: **proxy_ssl_certificate**, **proxy_ssl_certificate_key**, **grpc_ssl_certificate**, **grpc_ssl_certificate_key**, **uwsgi_ssl_certificate**, and **uwsgi_ssl_certificate_key**.
- The **listen** directive in the stream module now supports a new **fastopen** parameter, which enables **TCP Fast Open** mode for listening sockets.
- A new **max_errors** directive has been added to the **mail** proxy module.

Other changes:

- **nginx** now always returns an error if:
 - The **CONNECT** method is used.
 - Both **Content-Length** and **Transfer-Encoding** headers are specified in the request.
 - The request header name contains spaces or control characters.
 - The **Host** request header line contains spaces or control characters.
- **nginx** now blocks all HTTP/1.0 requests that include the **Transfer-Encoding** header.
- **nginx** now establishes HTTP/2 connections using the Application Layer Protocol Negotiation (ALPN) and no longer supports the Next Protocol Negotiation (NPN) protocol.

To install the **nginx:1.22** stream, use:

```
# yum module install nginx:1.22
```

If you want to upgrade from the **nginx:1.20** stream, see [Switching to a later stream](#).

For more information, see [Setting up and configuring NGINX](#).

For information about the length of support for the **nginx** module streams, see the [Red Hat Enterprise Linux Application Streams Life Cycle](#).

Bugzilla:2112345

mod_security rebased to version 2.9.6

The **mod_security** module for the Apache HTTP Server has been updated to version 2.9.6, which provides new features, bug fixes, and security fixes over the previously available version 2.9.2.

Notable enhancements include:

- Adjusted parser activation rules in the **modsecurity.conf-recommended** file.
- Enhancements to the way **mod_security** parses HTTP multipart requests.
- Added a new **MULTIPART_PART_HEADERS** collection.
- Added microsec timestamp resolution to the formatted log timestamp.
- Added missing Geo Countries.

Bugzilla:2143207

New packages: tomcat

RHEL 8.8 introduces the Apache Tomcat server version 9. Tomcat is the servlet container that is used in the official Reference Implementation for the Java Servlet and JavaServer Pages technologies. The Java Servlet and JavaServer Pages specifications are developed by Sun under the Java Community Process. Tomcat is developed in an open and participatory environment and released under the Apache Software License version 2.0.

Bugzilla:2160455

A new module stream: postgresql:15

RHEL 8.8 introduces **PostgreSQL 15**, which provides a number of new features and enhancements over version 13. Notable changes include:

- You can now access **PostgreSQL** JSON data by using subscripts. Example query:


```
SELECT ('{"postgres": { "release": 15 }}'::jsonb)['postgres']['release'];
```
- **PostgreSQL** now supports multirange data types and extends the **range_agg** function to aggregate multirange data types.
- **PostgreSQL** improves monitoring and observability:
 - You can now track progress of the **COPY** commands and Write-ahead-log (WAL) activity.
 - **PostgreSQL** now provides statistics on replication slots.
 - By enabling the **compute_query_id** parameter, you can now uniquely track a query through several **PostgreSQL** features, including **pg_stat_activity** or **EXPLAIN VERBOSE**.
- **PostgreSQL** improves support for query parallelism by the following:

- Improved performance of parallel sequential scans.
- The ability of SQL Procedural Language (**PL/pgSQL**) to execute parallel queries when using the **RETURN QUERY** command.
- Enabled parallelism in the **REFRESH MATERIALIZED VIEW** command.
- **PostgreSQL** now includes the SQL standard **MERGE** command. You can use **MERGE** to write conditional SQL statements that can include the **INSERT**, **UPDATE**, and **DELETE** actions in a single statement.
- **PostgreSQL** provides the following new functions for using regular expressions to inspect strings: **regexp_count()**, **regexp_instr()**, **regexp_like()**, and **regexp_substr()**.
- **PostgreSQL** adds the **security_invoker** parameter, which you can use to query data with the permissions of the view caller, not the view creator. This helps you ensure that view callers have the correct permissions for working with the underlying data.
- **PostgreSQL** improves performance, namely in its archiving and backup facilities.
- **PostgreSQL** adds support for the **LZ4** and **Zstandard (zstd)** lossless compression algorithms.
- **PostgreSQL** improves its in-memory and on-disk sorting algorithms.
- The updated **postgresql.service** systemd unit file now ensures that the **postgresql** service is started after the network is up.

The following changes are backwards incompatible:

- The default permissions of the public schema have been modified. Newly created users need to grant permission explicitly by using the **GRANT ALL ON SCHEMA public TO myuser;** command. For example:

```
postgres=# CREATE USER mydbuser;
postgres=# GRANT ALL ON SCHEMA public TO mydbuser;
postgres=# \c postgres mydbuser
postgres=# CREATE TABLE mytable (id int);
```

- The **libpq PQsendQuery()** function is no longer supported in pipeline mode. Modify affected applications to use the **PQsendQueryParams()** function instead.

See also [Using PostgreSQL](#).

To install the **postgresql:15** stream, use:

```
# yum module install postgresql:15
```

If you want to upgrade from an earlier **postgresql** stream within RHEL 8, follow the procedure described in [Switching to a later stream](#) and then migrate your **PostgreSQL** data as described in [Migrating to a RHEL 8 version of PostgreSQL](#).

For information about the length of support for the **postgresql** module streams, see the [Red Hat Enterprise Linux Application Streams Life Cycle](#).

[Bugzilla:2128241](#)

4.11. COMPILERS AND DEVELOPMENT TOOLS

A new module stream: **swig:4.1**

RHEL 8.8 introduces the Simplified Wrapper and Interface Generator (SWIG) version 4.1, available as a new module stream, **swig:4.1**.

Compared to **SWIG 4.0** released in RHEL 8.4, **SWIG 4.1**:

- Adds support for **Node.js** versions 12 to 18 and removes support for **Node.js** versions earlier than 6.
- Adds support for **PHP 8**.
- Handles **PHP** wrapping entirely through **PHP** C API and no longer generates a **.php** wrapper by default.
- Supports only **Perl 5.8.0** and later versions.
- Adds support for **Python** versions 3.9 to 3.11.
- Supports only **Python 3.3** and later **Python 3** versions, and **Python 2.7**.
- Provides fixes for various memory leaks in **Python**-generated code.
- Improves support for the C99, C++11, C++14, and C++17 standards and starts implementing the C++20 standard.
- Adds support for the C++ **std::unique_ptr** pointer class.
- Includes several minor improvements in C++ template handling.
- Fixes C++ declaration usage in various cases.

To install the **swig:4.1** module stream, use:

```
# yum module install swig:4.1
```

If you want to upgrade from an earlier **swig** module stream, see [Switching to a later stream](#).

For information about the length of support for the **swig** module streams, see the [Red Hat Enterprise Linux Application Streams Life Cycle](#).

[Bugzilla:2139076](#)

A new module stream: **jaxb:4**

RHEL 8.8 introduces Jakarta XML Binding (JAXB) 4 as the new **jaxb:4** module stream. JAXB is a framework that enables developers to map Java classes to and from XML representations.

To install the **jaxb:4** module stream, use:

```
# yum module install jaxb:4
```

[Bugzilla:2055539](#)

Updated GCC Toolset 12

GCC Toolset 12 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.

Notable changes introduced in RHEL 8.8 include:

- The GCC compiler has been updated to version 12.2.1, which provides many bug fixes and enhancements that are available in upstream GCC.
- **annobin** has been updated to version 11.08.

The following tools and versions are provided by GCC Toolset 12:

Tool	Version
GCC	12.2.1
GDB	11.2
binutils	2.38
dwz	0.14
annobin	11.08

To install GCC Toolset 12, run the following command as root:

```
# yum install gcc-toolset-12
```

To run a tool from GCC Toolset 12:

```
$ scl enable gcc-toolset-12 tool
```

To run a shell session where tool versions from GCC Toolset 12 override system versions of these tools:

```
$ scl enable gcc-toolset-12 bash
```

For more information, see [GCC Toolset 12](#).

[Bugzilla:2110582](#)

Security improvements added for **glibc**

The **SafeLinking** feature has been added to **glibc**. As a result, it improves protection for the **malloc** family of functions against certain single-linked list corruption including the allocator's thread-local cache.

[Bugzilla:1871383](#)

Improved **glibc** dynamic loader algorithm

The **glibc** dynamic loader's $O(n^3)$ algorithm for processing shared objects could result in slower application startup and shutdown times when shared object dependencies are deeply nested. With this

update, the dynamic loader’s algorithm has been improved to use a depth-first search (DFS). As a result, application startup and shutdown times are greatly improved in cases where shared object dependencies are deeply nested.

You can select the dynamic loader’s $O(n^3)$ algorithm by using the **glibc** runtime tunable **glibc.rtdl.dynamic_sort**. The default value of the tunable is 2, representing the new DFS algorithm. To select the previous $O(n^3)$ algorithm for compatibility, set the tunable to 1:

```
# GLIBC_TUNABLES=glibc.rtdl.dynamic_sort=1
# export GLIBC_TUNABLES
```

Bugzilla:1159809

LLVM Toolset rebased to version 15.0.7

LLVM Toolset has been updated to version 15.0.7. Notable changes include:

- The **-Wimplicit-function-declaration** and **-Wimplicit-int** warnings are enabled by default in C99 and newer. These warnings will become errors by default in Clang 16 and beyond.

Bugzilla:2118568

Rust Toolset rebased to version 1.66.1

Rust Toolset has been updated to version 1.66.1. Notable changes include:

- The **thread::scope** API creates a lexical scope in which local variables can be safely borrowed by newly spawned threads, and those threads are all guaranteed to exit before the scope ends.
- The **hint::black_box** API adds a barrier to compiler optimization, which is useful for preserving behavior in benchmarks that might otherwise be optimized away.
- The **.await** keyword now makes conversions with the **IntoFuture** trait, similar to the relationship between **for** and **Intolterator**.
- Generic associated types (GATs) allow traits to include type aliases with generic parameters, enabling new abstractions over both types and lifetimes.
- A new **let-else** statement allows binding local variables with conditional pattern matching, executing a divergent **else** block when the pattern does not match.
- Labeled blocks allow **break** statements to jump to the end of the block, optionally including an expression value.
- **rust-analyzer** is a new implementation of the Language Server Protocol, enabling Rust support in many editors. This replaces the former **rls** package, but you might need to adjust your editor configuration to migrate to **rust-analyzer**.
- Cargo has a new **cargo remove** subcommand for removing dependencies from **Cargo.toml**.

Bugzilla:2123899

Go Toolset rebased to version 1.19.4

Go Toolset has been updated to version 1.19.4. Notable changes include:

- Security fixes to the following packages:

- **crypto/tls**
- **mime/multipart**
- **net/http**
- **path/filepath**
- Bug fixes to:
 - The **go** command
 - The linker
 - The runtime
 - The **crypto/x509** package
 - The **net/http** package
 - The **time** package

Bugzilla:2174430

The **tzdata** package now includes the **/usr/share/zoneinfo/leap-seconds.list** file

Previously, the **tzdata** package only shipped the **/usr/share/zoneinfo/leapseconds** file. Some applications rely on the alternate format provided by the **/usr/share/zoneinfo/leap-seconds.list** file and, as a consequence, would experience errors.

With this update, the **tzdata** package now includes both files, supporting applications that rely on either format.

[Bugzilla:2154109](#)

4.12. IDENTITY MANAGEMENT

SSSD support for converting home directories to lowercase

With this enhancement, you can now configure SSSD to convert user home directories to lowercase. This helps to integrate better with the case-sensitive nature of the RHEL environment. The **override_homedir** option in the **[nss]** section of the **/etc/sss/sss.conf** file now recognizes the **%h** template value. If you use **%h** as part of the **override_homedir** definition, SSSD replaces **%h** with the user's home directory in lowercase.

Jira:RHELPLAN-139430

The **ipapwpolicy ansible-freeipa** module now supports new password policy options

With this update, the **ipapwpolicy** module included in the **ansible-freeipa** package supports additional **libpwquality** library options:

maxrepeat

Specifies the maximum number of the same character in sequence.

maxsequence

Specifies the maximum length of monotonic character sequences (**abcd**).

dictcheck

Checks if the password is a dictionary word.

usercheck

Checks if the password contains the username.

If any of the new password policy options are set, the minimum length of passwords is 6 characters. The new password policy settings are applied only to new passwords.

In a mixed environment with RHEL 7 and RHEL 8 servers, the new password policy settings are enforced only on servers running on RHEL 8.4 and later. If a user is logged in to an IdM client and the IdM client is communicating with an IdM server running on RHEL 8.3 or earlier, then the new password policy requirements set by the system administrator do not apply. To ensure consistent behavior, upgrade all servers to RHEL 8.4 and later.

Jira:RHELPLAN-137416

IdM now supports the ipanetgroup Ansible management module

As an Identity Management (IdM) system administrator, you can integrate IdM with NIS domains and netgroups. Using the **ipanetgroup ansible-freeipa** module, you can achieve the following:


- You can ensure that an existing IdM netgroup contains specific IdM users, groups, hosts and host groups and nested IdM netgroups.
- You can ensure that specific IdM users, groups, hosts and host groups and nested IdM netgroups are absent from an existing IdM netgroup.
- You can ensure that a specific netgroup is present or absent in IdM.

Jira:RHELPLAN-137411

New ipaclient_configure_dns_resolver and ipaclient_dns_servers Ansible ipaclient role variables specifying the client's DNS resolver

Previously, when using the **ansible-freeipa ipaclient** role to install an Identity Management (IdM) client, it was not possible to specify the DNS resolver during the installation process. You had to configure the DNS resolver before the installation.

With this enhancement, you can specify the DNS resolver when using the **ipaclient** role to install an IdM client with the **ipaclient_configure_dns_resolver** and **ipaclient_dns_servers** variables. Consequently, the **ipaclient** role modifies the **resolv.conf** file and the **NetworkManager** and **systemd-resolved** utilities to configure the DNS resolver on the client in a similar way that the **ansible-freeipa ipaserver** role does on the IdM server. As a result, configuring DNS when using the **ipaclient** role to install an IdM client is now more efficient.

**NOTE**

Using the **ipa-client-install** command-line installer to install an IdM client still requires configuring the DNS resolver before the installation.

Jira:RHELPLAN-137406

Using the ipaclient role to install an IdM client with an OTP requires no prior modification of the Ansible controller

Previously, the **kinit** command on the Ansible controller was a prerequisite for obtaining a one-time-

password (OTP) for Identity Management (IdM) client deployment. The need to obtain the OTP on the controller was a problem for Red Hat Ansible Automation Platform (AAP), where the **krb5-workstation** package was not installed by default.

With this update, the request for the administrator's TGT is now delegated to the first specified or discovered IdM server. As a result, you can now use an OTP to authorize the installation of an IdM client with no additional modification of the Ansible controller. This simplifies using the **ipaclient** role with AAP.

Jira:RHELPLAN-137403

SSSD now supports changing LDAP user passwords with the **shadow** password policy

With this enhancement, if you set **ldap_pwd_policy** to **shadow** in the `/etc/sss/sss.conf` file, LDAP users can now change their password stored in LDAP. Previously, password changes were rejected if **ldap_pwd_policy** was set to **shadow** as it was not clear if the corresponding **shadow** LDAP attributes were being updated.

Additionally, if the LDAP server cannot update the **shadow** attributes automatically, set the **ldap_chpass_update_last_change** option to **True** in the `/etc/sss/sss.conf` file to indicate to SSSD to update the attribute.

Bugzilla:2144519

Configure **pam_pwhistory** using a configuration file

With this update, you can configure the **pam_pwhistory** module in the `/etc/security/pwhistory.conf` configuration file. The **pam_pwhistory** module saves the last password for each user in order to manage password change history. Support has also been added in **authselect** which allows you to add the **pam_pwhistory** module to the PAM stack.

[Bugzilla:2068461](#), [Bugzilla:2063379](#)

getcert add-scep-ca now checks if user-provided SCEP CA certificates are in a valid PEM format

To add a SCEP CA to **certmonger** using the **getcert add-scep-ca** command, the provided certificate must be in a valid PEM format. Previously, the command did not check the user-provided certificate and did not return an error in case of an incorrect format. With this update, **getcert add-scep-ca** now checks the user-provided certificate and returns an error if the certificate is not in the valid PEM format.

[Bugzilla:2150025](#)

IdM now supports new Active Directory certificate mapping templates

Active Directory (AD) domain administrators can manually map certificates to a user in AD using the **altSecurityIdentities** attribute. There are six supported values for this attribute, though three mappings are now considered insecure. As part of [May 10, 2022 security update](#), once this update is installed on a domain controller, all devices are in compatibility mode. If a certificate is weakly mapped to a user, authentication occurs as expected but a warning message is logged identifying the certificates that are not compatible with full enforcement mode. As of November 14, 2023 or later, all devices will be updated to full enforcement mode and if a certificate fails the strong mapping criteria, authentication will be denied.

IdM now supports the new mapping templates, making it easier for an AD administrator to use the new rules and not maintain both. IdM now supports the following new mapping templates :

- Serial Number: **LDAPU1:(altSecurityIdentities=X509:<I>{issuer_dn!ad_x500}<SR>{serial_number!hex_ur})**

- Subject Key Id: **LDAPU1:(altSecurityIdentities=X509:<SKI>{subject_key_id!hex_u})**
- User SID: **LDAPU1:(objectsid={sid})**

If you do not want to reissue certificates with the new SID extension, you can create a manual mapping by adding the appropriate mapping string to a user's **altSecurityIdentities** attribute in AD.

[Bugzilla:2087247](#)

samba rebased to version 4.17.5

The **samba** packages have been upgraded to upstream version 4.17.5, which provides bug fixes and enhancements over the previous version. The most notable changes:

- Security improvements in previous releases impacted the performance of the Server Message Block (SMB) server for high meta data workloads. This update improves the performance in this scenario.
- The **--json** option was added to the **smbstatus** utility to display detailed status information in JSON format.
- The **samba.smb.conf** and **samba.samba3.smb.conf** modules have been added to the **smbconf** Python API. You can use them in Python programs to read and, optionally, write the Samba configuration natively.

Note that the server message block version 1 (SMB1) protocol is deprecated since Samba 4.11 and will be removed in a future release.

Back up the database files before starting Samba. When the **smbd**, **nmbd**, or **winbind** services start, Samba automatically updates its **tdb** database files. Red Hat does not support downgrading **tdb** database files.

After updating Samba, use the **testparm** utility to verify the **/etc/samba/smb.conf** file.

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

[Bugzilla:2132051](#)

ipa-client-install now supports authentication with PKINIT

Previously, the **ipa-client-install** supported only password based authentication. This update provides support to **ipa-client-install** for authentication with PKINIT.

For example:

```
ipa-client-install --pkinit-identity=FILE:/path/to/cert.pem,/path/to/key.pem --pkinit-
anchor=FILE:/path/to/cacerts.pem
```

To use the PKINIT authentication, you must establish trust between IdM and the CA chain of the PKINIT certificate. For more information see the **ipa-cacert-manage(1)** man page. Also, the certificate identity mapping rules must map the PKINIT certificate of the host to a principal that has permission to add or modify a host record. For more information see the **ipa certmaprule-add** man page.

[Bugzilla:2075452](#)

Directory server now supports ECDSA private keys for TLS

Previously, you could not use cryptographic algorithms that are stronger than RSA to secure Directory Server connections. With this enhancement, Directory Server now supports both ECDSA and RSA keys.

[Bugzilla:2096795](#)

New `pamModuleIsThreadSafe` configuration option is now available

When a PAM module is thread-safe, you can improve the PAM authentication throughput and response time of that specific module, by setting the new `pamModuleIsThreadSafe` configuration option to **yes**:

```
`pamModuleIsThreadSafe: yes`
```

This configuration applies on the PAM module configuration entry (child of **cn=PAM Pass Through Auth,cn=plugins,cn=config**).

Use `pamModuleIsThreadSafe` option in the `dse.ldif` configuration file or the `ldapmodify` command. Note that the `ldapmodify` command requires you to restart the server.

[Bugzilla:2142639](#)

New `nsslapd-auditlog-display-attrs` configuration parameter for the Directory Server audit log

Previously, the distinguished name (DN) was the only way to identify the target entry in the audit log event. With the new `nsslapd-auditlog-display-attrs` parameter, you can configure Directory Server to display additional attributes in the audit log, providing more details about the modified entry..

For example, if you set the `nsslapd-auditlog-display-attrs` parameter to **cn**, the audit log displays the entry **cn** attribute in the output. To include all attributes of a modified entry, use an asterisk (`*`) as the parameter value.

For more information, see [nsslapd-auditlog-display-attrs](#).

[Bugzilla:2136610](#)

4.13. DESKTOP

The `inkscape1` package replaces `inkscape`

With this release, the new, non-modular `inkscape1` package replaces the legacy, modular `inkscape` package. This also upgrades the Inkscape application from version 0.92 to version 1.0.

Inkscape 1.0 no longer depends on the Python 2 runtime and instead uses Python 3.

For the complete list of changes in Inkscape 1.0, see the upstream release notes: <https://inkscape.org/release/inkscape-1.0/>.

Jira:RHELPLAN-121672

Kiosk mode supports an on-screen keyboard

You can now use the GNOME on-screen keyboard (OSK) in the kiosk mode session.

To enable the OSK, select the **Kiosk (with on-screen keyboard)** option from the gear menu at the login screen.

Note that kiosk mode in RHEL 8 is based on the X11 protocol, which causes certain known issues with the OSK. Notably, you cannot type accented characters, such as **é** or **ü**, on the OSK. See [BZ#1916470](#) for details.

[Bugzilla:2070976](#)

Support for NTLMv2 in **libsoup** and Evolution

The **libsoup** library can now authenticate with the Microsoft Exchange Server using the NT LAN Manager version 2 (NTLMv2) protocol. Previously, **libsoup** supported only the NTLMv1 protocol, which might be disabled in certain configurations due to security issues.

As a result, Evolution and other applications that internally use **libsoup** can also authenticate with the Microsoft Exchange Server using NTLMv2.

[Bugzilla:1938011](#)

Custom right-click menu on the desktop

You can now customize the menu that opens when you right-click the desktop background. You can create custom entries in the menu that run arbitrary commands.

To customize the menu, see [Customizing the right-click menu on the desktop](#).

[Bugzilla:2033572](#)

Disable swipe to switch workspaces

Previously, swiping up or down with three fingers always switched the workspace on a touch screen. With this release, you can disable the workspace switching.

For details, see [Disabling swipe to switch workspaces](#).

[Bugzilla:2138109](#)

4.14. THE WEB CONSOLE

The web console now performs additional steps for binding LUKS-encrypted root volumes to NBDE

With this update, the RHEL web console performs additional steps required for binding LUKS-encrypted root volumes to Network-Bound Disk Encryption (NBDE) deployments. After you select an encrypted root file system and a Tang server, you can skip adding the **rd.neednet=1** parameter to the kernel command line, installing the **clevis-dracut** package, and regenerating an initial ramdisk (**initrd**). For non-root file systems, the web console now enables the **remote-cryptsetup.target** and **clevis-luks-akspass.path systemd** units, installs the **clevis-systemd** package, and adds the **_netdev** parameter to the **fstab** and **crypttab** configuration files. As a result, you can now use the graphical interface for all Clevis-client configuration steps when creating NBDE deployments for automated unlocking of LUKS-encrypted root volumes.

Jira:RHELPLAN-139125

Certain cryptographic subpolicies are now available in the web console

This update of the RHEL web console extends the options in the **Change crypto policy** dialog. Besides the four system-wide cryptographic policies, you can also apply the following subpolicies through the graphical interface now:


- **DEFAULT:SHA1** is the **DEFAULT** policy with the **SHA-1** algorithm enabled.
- **LEGACY:AD-SUPPORT** is the **LEGACY** policy with less secure settings that improve interoperability for Active Directory services.
- **FIPS:OSPP** is the **FIPS** policy with further restrictions inspired by the Common Criteria for Information Technology Security Evaluation standard.

Jira:RHELPLAN-137505

4.15. RED HAT ENTERPRISE LINUX SYSTEM ROLES

New IPsec customization parameters for the **vpn** RHEL System Role

Because certain network devices require IPsec customization to work correctly, the following parameters have been added to the **vpn** RHEL System Role:



IMPORTANT

Do not change the following parameters without advanced knowledge. Most scenarios do not require their customization.

Furthermore, for security reasons, encrypt a value of the **shared_key_content** parameter by using Ansible Vault.

- Tunnel parameters:
 - **shared_key_content**
 - **ike**
 - **esp**
 - **ikelifetime**
 - **salifetime**
 - **retransmit_timeout**
 - **dpddelay**
 - **dpdtimeout**
 - **dpdaction**
 - **leftupdown**
- Per-host parameters:
 - **leftid**
 - **rightid**

As a result, you can use the **vpn** role to configure IPsec connectivity to a wide range of network devices.

[Bugzilla:2119600](#)

The `ha_cluster` System Role now supports automated execution of the `firewall`, `selinux`, and `certificate` System Roles

The `ha_cluster` RHEL System Role now supports the following features:

Using the `firewall` and `selinux` System Roles to manage port access

To configure the ports of a cluster to run the `firewalld` and `selinux` services, you can set the new role variables `ha_cluster_manage_firewall` and `ha_cluster_manage_selinux` to `true`. This configures the cluster to use the `firewall` and `selinux` System Roles, automating and performing these operations within the `ha_cluster` System Role. If these variables are set to their default value of `false`, the roles are not performed. With this release, the firewall is no longer configured by default, because it is configured only when `ha_cluster_manage_firewall` is set to `true`.

Using the `certificate` System Role to create `apcsd` private key and certificate pair

The `ha_cluster` System Role now supports the `ha_cluster_pcsd_certificates` role variable. Setting this variable passes on its value to the `certificate_requests` variable of the `certificate` System Role. This provides an alternative method for creating the private key and certificate pair for `pcs`.

[Bugzilla:2130019](#)

The `ha_cluster` System Role now supports quorum device configuration

A quorum device acts as a third-party arbitration device for a cluster. A quorum device is recommended for clusters with an even number of nodes. With two-node clusters, the use of a quorum device can better determine which node survives in a split-brain situation. You can now configure a quorum device with the `ha_cluster` System Role, both `qdevice` for a cluster and `qnetd` for an arbitration node.

[Bugzilla:2143814](#)

The `metrics` System Role does not work with disabled fact gathering

Ansible fact gathering might be disabled in your environment for performance or other reasons. In such configurations, it is not currently possible to use the `metrics` System Role. To work around this problem, enable fact caching, or do not use the `metrics` System Role if it is not possible to use fact gathering.

[Bugzilla:2079009](#)

The `postfix` RHEL System Role can now use the `firewall` and `selinux` RHEL System Roles to manage port access

With this enhancement, you can automate managing port access by using the new role variables `postfix_manage_firewall` and `postfix_manage_selinux`:

- If they are set to `true`, each role is used to manage the port access.
- If they are set to `false`, which is default, the roles do not engage.

[Bugzilla:2130332](#)

The `vpn` RHEL System Role can now use the `firewall` and `selinux` roles to manage port access

With this enhancement, you can automate managing port access in the `vpn` RHEL System Role through the `firewall` and `selinux` roles. If you set the new role variables `vpn_manage_firewall` and `vpn_manage_selinux` to `true`, the roles manage port access.

[Bugzilla:2130345](#)

The **metrics** RHEL System Role now can use the **firewall** role and the **selinux** role to manage port access

With this enhancement, you can control access to ports. If you set the new role variables **metrics_manage_firewall** and **metrics_manage_selinux** to **true**, the roles will manage port access. You can now automate and perform these operations directly by using the **metrics** role.

[Bugzilla:2133532](#)

The **nbde_server** RHEL System Role now can use the **firewall** and **selinux** roles to manage port access

With this enhancement, you can use the **firewall** and **selinux** roles to manage ports access. If you set the new role variables **nbde_server_manage_firewall** and **nbde_server_manage_selinux** to **true**, the roles manage port access. You can now automate these operations directly by using the **nbde_server** role.

[Bugzilla:2133931](#)

The **initscripts** network provider supports route metric configuration of the default gateway

With this update, you can use the **initscripts** network provider in the **rhel-system-roles.network** RHEL System Role to configure the route metric of the default gateway.

The reasons for such a configuration could be:

- Distributing the traffic load across the different paths
- Specifying primary routes and backup routes
- Leveraging routing policies to send traffic to specific destinations through specific paths

[Bugzilla:2134201](#)

The **network** System Role supports setting a DNS priority value

This enhancement adds the **dns_priority** parameter to the RHEL **network** System Role. You can set this parameter to a value from **-2147483648** to **2147483647**. The default value is **0**. Lower values have a higher priority. Note that negative values cause the System Role to exclude other configurations with a greater numeric priority value. Consequently, in presence of at least one negative priority value, the System Role uses only DNS servers from connection profiles with the lowest priority value.

As a result, you can use the **network** System Role to define the order of DNS servers in different connection profiles.

[Bugzilla:2133856](#)

Added support for the cloned MAC address

Cloned MAC address is the MAC address of the device WAN port which is the same as the MAC address of the machine. With this update, users can specify the bonding or bridge interface with the MAC address or the strategy such as **random** or **preserve** to get the default MAC address for the bonding or bridge interface.

[Bugzilla:2143458](#)

The **cockpit** RHEL System Role integration with the **firewall**, **selinux**, and **certificate** roles

This enhancement enables you to integrate the **cockpit** role with the **firewall** role and the **selinux** role to manage port access and the **certificate** role to generate certificates.

To control the port access, use the new **cockpit_manage_firewall** and **cockpit_manage_selinux** variables. Both variables are set to **false** by default and are not executed. Set them to **true** to allow the **firewall** and **selinux** roles to manage the RHEL web console service port access. The operations will then be executed within the **cockpit** role.

Note that you are responsible for managing port access for firewall and SELinux.

To generate certificates, use the new **cockpit_certificates** variable. The variable is set to **false** by default and is not executed. You can use this variable the same way you would use the **certificate_request** variable in the **certificate** role. The **cockpit** role will then use the **certificate** role to manage the RHEL web console certificates.

[Bugzilla:2137667](#)

The **selinux** RHEL System Role now supports the **local** parameter

This update of the **selinux** RHEL System Role introduces support for the **local** parameter. By using this parameter, you can remove only your local policy modifications and preserve the built-in SELinux policy.

[Bugzilla:2143385](#)

New RHEL System Role for direct integration with Active Directory

The new **rhel-system-roles.ad_integration** RHEL System Role was added to the **rhel-system-roles** package. As a result, administrators can now automate direct integration of a RHEL system with an Active Directory domain.

[Bugzilla:2144876](#)

New Ansible Role for Red Hat Insights and subscription management

The **rhel-system-roles** package now includes the remote host configuration (**rhc**) system role. This role enables administrators to easily register RHEL systems to Red Hat Subscription Management (RHSM) and Satellite servers. By default, when you register a system by using the **rhc** system role, the system connects to Red Hat Insights. With the new **rhc** system role, administrators can now automate the following tasks on the managed nodes:

- Configure the connection to Red Hat Insights, including automatic update, remediations, and tags for the system.
- Enable and disable repositories.
- Configure the proxy to use for the connection.
- Set the release of the system.

For more information about how to automate these tasks, see [Using the RHC system role to register the system](#).

[Bugzilla:2144877](#)

Microsoft SQL Server Ansible role supports asynchronous high availability replicas

Previously, Microsoft SQL Server Ansible role supported only primary, synchronous, and witness high availability replicas. Now, you can set the **mssql_ha_replica_type** variable to **asynchronous** to configure it with asynchronous replica type for a new or existing replica.

[Bugzilla:2144820](#)

Microsoft SQL Server Ansible role supports the read-scale cluster type

Previously, Microsoft SQL Ansible role supported only the external cluster type. Now, you can configure the role with a new variable **mssql_ha_ag_cluster_type**. The default value is **external**, use it to configure the cluster with Pacemaker. To configure the cluster without Pacemaker, use the value **none** for that variable.

[Bugzilla:2144821](#)

Microsoft SQL Server Ansible role can generate TLS certificates

Previously, you needed to generate a TLS certificate and a private key on the nodes manually before configuring the Microsoft SQL Ansible role. With this update, the Microsoft SQL Server Ansible role can use the **redhat.rhel_system_roles.certificate** role for that purpose. Now, you can set the **mssql_tls_certificates** variable in the format of the **certificate_requests** variable of the **certificate** role to generate a TLS certificate and a private key on the node.

[Bugzilla:2144852](#)

Microsoft SQL Server Ansible role supports configuring SQL Server version 2022

Previously, Microsoft SQL Ansible role supported only configuring SQL Server version 2017 and version 2019. This update provides you with the support for SQL Server version 2022 for Microsoft SQL Ansible role. Now, you can set **mssql_version** value to **2022** for configuring a new SQL Server 2022 or upgrading SQL Server from version 2019 to version 2022. Note that upgrade of an SQL Server from version 2017 to version 2022 is unavailable.

[Bugzilla:2153428](#)

Microsoft SQL Server Ansible role supports configuration of the Active Directory authentication

With this update, the Microsoft SQL Ansible role supports configuration of the Active Directory authentication for an SQL Server. Now, you can configure the Active Directory authentication by setting variables with the **mssql_ad_** prefix.

[Bugzilla:2163696](#)

The logging RHEL System Role integration with the firewall, selinux, and certificate roles

This enhancement enables you to integrate the **logging** role with the **firewall** role and the **selinux** role to manage port access and the **certificate** role to generate certificates.

To control the port access, use the new **logging_manage_firewall** and **logging_manage_selinux** variables. Both variables are set to **false** by default and are not executed. Set them to **true** to execute the roles within the **logging** role.

Note that you are responsible for managing port access for firewall and SELinux.

To generate certificates, use the new **logging_certificates** variable. The variable is set to **false** by default and the **certificate** role is not executed. You can use this variable the same way you would use the **certificate_request** variable in the **certificate** role. The **logging** role will then use the **certificate** role to manage the certificates.

[Bugzilla:2130362](#)

Routing rule is able to look up a route table by its name

With this update, the **rhel-system-roles.network** RHEL System Role supports looking up a route table by its name when you define a routing rule. This feature provides quick navigation for complex network configurations where you need to have different routing rules for different network segments.

[Bugzilla:2129620](#)

Microsoft SQL Server Ansible role supports configuring SQL Server version 2022

Previously, Microsoft SQL Ansible role supported only configuring SQL Server version 2017 and version 2019. This update provides you with the support for SQL Server version 2022 for Microsoft SQL Ansible role. Now, you can set **mssql_version** value to **2022** for configuring a new SQL Server 2022 or upgrading SQL Server from version 2019 to version 2022. Note that upgrade of an SQL Server from version 2017 to version 2022 is unavailable.

[Bugzilla:2153427](#)

The **journald** RHEL System Role is now available

The **journald** service collects and stores log data in a centralized database. With this enhancement, you can use the **journald** System Role variables to automate the configuration of the **systemd** journal, and configure persistent logging by using the Red Hat Ansible Automation Platform.

[Bugzilla:2165176](#)

The **sshd** RHEL System Role can now use the **firewall** and **selinux** RHEL System Roles to manage port access

With this enhancement, you can automate managing port access by using the new role variables **sshd_manage_firewall** and **sshd_manage_selinux**. If they are set to **true**, each role is used to manage the port access. If they are set to **false**, which is default, the roles do not engage.

[Bugzilla:2149683](#)

4.16. VIRTUALIZATION

Hardware cryptographic devices can now be automatically hot-plugged

Previously, it was only possible to define cryptographic devices for passthrough if they were present on the host before the mediated device was started. Now, you can define a mediated device matrix that lists all the cryptographic devices that you want to pass through to your virtual machine (VM). As a result, the specified cryptographic devices are automatically passed through to the running VM if they become available later. Also, if the devices become unavailable, they are removed from the VM, but the guest operating system keeps running normally.

[Bugzilla:1660908](#)

Improved performance for PCI passthrough devices on IBM Z

With this update, the PCI passthrough implementation on IBM Z hardware has been enhanced through multiple improvements to I/O handling. As a result, PCI devices passed through to KVM virtual machines (VMs) on IBM Z hosts now have significantly better performance.

In addition, ISM devices can now be assigned to VMs on IBM Z hosts.

[Bugzilla:1664379](#)

RHEL 8 guests now support SEV-SNP

On virtual machines (VMs) that use RHEL 8 as a guest operating system, you can now use AMD Secure Encrypted Virtualization (SEV) with the Secure Nested Paging (SNP) feature. Among other benefits, SNP enhances SEV by improving its memory integrity protection, which helps prevent hypervisor-based attacks such as data replay or memory re-mapping. Note that for SEV-SNP to work on a RHEL 8 VM, the host running the VM must support SEV-SNP as well.

Bugzilla:2087262

zPCI device assignment

It is now possible to attach zPCI devices as pass-through devices to virtual machines (VMs) hosted on RHEL running on IBM Z hardware. For example, this enables the use of NVMe flash drives in VMs.

Jira:RHELPLAN-59528

4.17. SUPPORTABILITY

The `sos` utility is moving to a 4-week update cadence

Instead of releasing **sos** updates with RHEL minor releases, the **sos** utility release cadence is changing from 6 months to 4 weeks. You can find details about the updates for the **sos** package in the RPM changelog every 4 weeks or you can read a summary of **sos** updates in the RHEL Release Notes every 6 months.

[Bugzilla:2164987](#)

The `sos clean` command now obfuscates IPv6 addresses

Previously, the **sos clean** command did not obfuscate IPv6 addresses, leaving some customer-sensitive data in the collected **sos** report. With this update, **sos clean** detects and obfuscates IPv6 addresses as expected.

[Bugzilla:2134906](#)

4.18. CONTAINERS

New `podman` RHEL System Role is now available

Beginning with Podman 4.2, you can use the **podman** System Role to manage Podman configuration, containers, and **systemd** services that run Podman containers.

Jira:RHELPLAN-118698

Podman now supports events for auditing

Beginning with Podman v4.4, you can gather all relevant information about a container directly from a single event and **journald** entry. To enable Podman auditing, modify the **container.conf** configuration file and add the **events_container_create_inspect_data=true** option to the **[engine]** section. The data is in JSON format, the same as from the **podman container inspect** command. For more information, see [How to use new container events and auditing features in Podman 4.4](#) .

Jira:RHELPLAN-136601

The Container Tools packages have been updated

The updated Container Tools packages, which contain the Podman, Buildah, Skopeo, crun, and runc tools, are now available. This update applies a series of bug fixes and enhancements over the previous version.

Notable changes in Podman v4.4 include:

- Introduce Quadlet, a new systemd-generator that easily creates and maintains systemd services using Podman.
- A new command, **podman network update**, has been added, which updates networks for containers and pods.
- A new command, **podman buildx version**, has been added, which shows the buildah version.
- Containers can now have startup healthchecks, allowing a command to be run to ensure the container is fully started before the regular healthcheck is activated.
- Support a custom DNS server selection using the **podman --dns** command.
- Creating and verifying sigstore signatures using Fulcio and Rekor is now available.
- Improved compatibility with Docker (new options and aliases).
- Improved Podman's Kubernetes integration - the commands **podman kube generate** and **podman kube play** are now available and replace the **podman generate kube** and **podman play kube** commands. The **podman generate kube** and **podman play kube** commands are still available but it is recommended to use the new **podman kube** commands.
- Systemd-managed pods created by the **podman kube play** command now integrate with sd-notify, using the **io.containers.sdnotify** annotation (or **io.containers.sdnotify/\$name** for specific containers).
- Systemd-managed pods created by **podman kube play** can now be auto-updated, using the **io.containers.auto-update** annotation (or **io.containers.auto-update/\$name** for specific containers).

Podman has been upgraded to version 4.4, for further information about notable changes, see [upstream release notes](#).

Jira:RHELPLAN-136608

Aardvark and Netavark now support custom DNS server selection

The Aardvark and Netavark network stack now support custom DNS server selection for containers instead of the default DNS servers on the host. You have two options for specifying the custom DNS server:

- Add the **dns_servers** field in the **containers.conf** configuration file.
- Use the new **--dns** Podman option to specify an IP address of the DNS server.

The **--dns** option overrides the values in the **container.conf** file.

Jira:RHELPLAN-138025

Skopeo now supports generating sigstore key pairs

You can use the **skopeo generate-sigstore-key** command to generate a sigstore public/private key pair. For more information, see **skopeo-generate-sigstore-key** man page.

Jira:RHELPLAN-151481

Toolbox is now available

With the **toolbox** utility, you can use the containerized command-line environment without installing troubleshooting tools directly on your system. Toolbox is built on top of Podman and other standard container technologies from OCI. For more information, see [toolbox](#).

Jira:RHELPLAN-150266

The capability for multiple trusted GPG keys for signing images is available

The **/etc/containers/policy.json** file supports a new **keyPaths** field which accepts a list of files containing the trusted keys. Because of this, the container images signed with Red Hat's General Availability and Beta GPG keys are now accepted in the default configuration.

For example:

```
"registry.redhat.io": [
  {
    "type": "signedBy",
    "keyType": "GPGKeys",
    "keyPaths": ["/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release", "/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-beta"]
  }
]
```

Jira:RHELPLAN-118470

RHEL 8 Extended Update Support

The RHEL Container Tools are now supported in RHEL 8 Extended Update Support (EUS) releases. More information on Red Hat Enterprise Linux EUS is available in [Container Tools AppStream - Content Availability](#), [Red Hat Enterprise Linux \(RHEL\) Extended Update Support \(EUS\) Overview](#) .

Jira:RHELPLAN-151121

The sigstore signatures are now available

Beginning with Podman 4.2, you can use the sigstore format of container image signatures. The sigstore signatures are stored in the container registry together with the container image without the need to have a separate signature server to store image signatures.

Jira:RHELPLAN-75165

Podman now supports the pre-execution hooks

The root-owned plugin scripts located in the **/usr/libexec/podman/pre-exec-hooks** and **/etc/containers/pre-exec-hooks** directories define a fine-control over container operations, especially blocking unauthorized actions.

The **/etc/containers/podman_preexec_hooks.txt** file must be created by an administrator and can be empty. If **/etc/containers/podman_preexec_hooks.txt** does not exist, the plugin scripts will not be executed. If all plugin scripts return zero value, then the **podman** command is executed, otherwise, the **podman** command exits with the inherited exit code.

Red Hat recommends using the following naming convention to execute the scripts in the correct order: **DDD-plugin_name.lang**, for example **010-check-group.py**. Note that the plugin scripts are valid at the time of creation. Containers created before plugin scripts are not affected.

Bugzilla:2119200

CHAPTER 5. 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 8.8. 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

nomodeset

With this kernel parameter, you can disable kernel mode setting. DRM drivers will not perform display-mode changes or accelerated rendering. Only the system frame buffer will be available for use if this was set-up by the firmware or boot loader.

nomodeset is useful as fallback, or for testing and debugging.

sev=option[,option...] [X86-64]

For more information, see [Documentation/x86/x86_64/boot-options.rst](#).

amd_pstate=[X86]

- **disable**: Do not enable **amd_pstate** as the default scaling driver for the supported processors.
- **passive**: Use **amd_pstate** as a scaling driver. The driver requests a desired performance on this abstract scale and the power management firmware translates the requests into actual hardware states, such as core frequency, data fabric and memory clocks and so on.

retbleed=ibpb,nosmt

This parameter is similar to **ibpb** and is an alternative for systems which do not have STIBP. With this parameter you can disable SMT when STIBP is not available.

Updated kernel parameters

amd_iommu=[HW,X86-64]

With this kernel parameter, you can pass parameters to the AMD IOMMU driver in the system. Possible values are:

- **fullflush**: Deprecated, equivalent to **iommu.strict=1**.
- **off**: do not initialize any AMD IOMMU found in the system.
- **force_isolation**: Force device isolation for all devices. The IOMMU driver is not allowed anymore to lift isolation requirements as needed.
 - This option does not override **iommu=pt**.
- **force_enable**: Force enable the IOMMU on platforms known to be buggy with IOMMU enabled.
 - Use this option with care.

crashkernel=size[KMG][@offset[KMG]]

[KNL] Using **kexec**, Linux can switch to a crash kernel upon panic. This parameter reserves the physical memory region [offset, offset + size] for that kernel image. If **@offset** is omitted, then a suitable offset is selected automatically.

[KNL, X86-64, ARM64] Select a region under 4G first, and fall back to reserve region above 4G when **@offset** has not been specified.

For more details, see [Documentation/admin-guide/kdump/kdump.rst](#).

crashkernel=size[KMG],low

- [KNL, X86-64, ARM64] With this parameter, you can specify low range under 4G for the second kernel. When **crashkernel=X,high** is passed, that require some amount of low memory, for example **swiotlb** requires at least 64M+32K low memory, also enough extra low memory is needed to make sure DMA buffers for 32-bit devices will not run out. Kernel would try to allocate default size of memory below 4G automatically. The default size is platform dependent.

- x86: `max(swiotlb_size_or_default() + 8MiB, 256MiB)`
- arm64: 128MiB
- 0: to disable low allocation.


This parameter will be ignored when **crashkernel=X,high** is not used or memory reserved is below 4G.

- [KNL, ARM64] With this parameter, you can specify a low range in the DMA zone for the crash dump kernel. This paramete will be ignored when **crashkernel=X,high** is not used.

intel_iommu=[DMAR]

The kernel parameter for setting the Intel IOMMU driver (DMAR) option.

- on: Enable intel iommu driver.
- off: Disable intel iommu driver.
- `igfx_off` [Default Off]: By default, gfx is mapped as normal device. If a gfx device has a dedicated DMAR unit, the DMAR unit is bypassed by not enabling DMAR with this option. In this case, the **gfx** device will use physical address for DMA.
- `strict` [Default Off]: Deprecated, equivalent to **`iommu.strict=1`**.
- `sp_off` [Default Off]: By default, super page will be supported if Intel IOMMU has the capability. With this option, super page will not be supported.
- `sm_on` [Default Off]: By default, scalable mode will be disabled even if the hardware advertises that it has support for the scalable mode translation. With this option set, scalable mode will be used on hardware which claims to support it.
- `tboot_noforce` [Default Off]: Do not force the Intel IOMMU enabled under **tboot**. By default, **tboot** will force Intel IOMMU on, which could harm performance of some high-throughput devices like 40Gbit network cards, even if identity mapping is enabled.

**NOTE**

Using this option lowers the security provided by **tboot** because it makes the system vulnerable to DMA attacks.

iommu.strict=[ARM64,X86]

With this kernel parameter, you can configure TLB invalidation behavior.

Format: { "0" | "1" }

- 0 - Lazy mode. Request that DMA unmap operations use deferred invalidation of hardware TLBs, for increased throughput at the cost of reduced device isolation. Will fall back to strict mode if not supported by the relevant IOMMU driver.
- 1 - Strict mode. DMA unmap operations invalidate IOMMU hardware TLBs synchronously.
- unset - Use value of **CONFIG_IOMMU_DEFAULT_DMA_{LAZY,STRICT}**.

**NOTE**

On x86, strict mode specified via one of the legacy driver-specific options takes precedence.

mem_encrypt=[X86-64]

The kernel parameter for setting the AMD Secure Memory Encryption (SME) control.

Valid arguments: on, off

Default depends on the kernel configuration option:

- on (CONFIG_AMD_MEM_ENCRYPT_ACTIVE_BY_DEFAULT=y)
- off (CONFIG_AMD_MEM_ENCRYPT_ACTIVE_BY_DEFAULT=n)
- mem_encrypt=on: Activate SME
- mem_encrypt=off: Do not activate SME
Refer to **Documentation/virt/kvm/x86/amd-memory-encryption.rst** for details on when memory encryption can be activated.

retbleed=[X86]

With this kernel parameter, you can control mitigation of RETBleed (Arbitrary Speculative Code Execution with Return Instructions) vulnerability.

AMD-based UNRET and IBPB mitigations alone do not stop sibling threads from influencing the predictions of other sibling threads. For that reason, STIBP is used on processors that support it, and mitigate SMT on processors that do not.

- off - no mitigation
- auto - automatically select a mitigation
- auto,nosmt - automatically select a mitigation, disabling SMT if necessary for the full mitigation (only on Zen1 and older without STIBP).
- ibpb - On AMD, mitigate short speculation windows on basic block boundaries too. Safe, highest performance impact. It also enables STIBP if present. Not suitable on Intel.

- `unret` - Force enable untrained return thunks, only effective on AMD f15h-f17h based systems.
- `unret,nosmt` - Like `unret`, but will disable SMT when STIBP is not available. This is the alternative for systems which do not have STIBP.

`swiotlb=[ARM,IA-64,PPC,MIPS,X86]`

With this kernel parameter, you can configure the behavior of I/O TLB slabs.

Format: { `<int>` [`,<int>`] | `force` | `noforce` }

- `<int>` - Number of I/O TLB slabs
- `<int>` - Second integer after comma. Number of `swiotlb` areas with their own lock. Must be power of 2.
- `force` - force using of bounce buffers even if they would not be automatically used by the kernel
- `noforce` - Never use bounce buffers (for debugging)

New `sysctl` parameters

`page_lock_unfairness`

This value determines the number of times that the page lock can be stolen from under a waiter. After the lock is stolen the number of times specified in this file (the default is 5), the *fair lock handoff* semantics will apply, and the waiter will only be awakened if the lock can be taken.

`rps_default_mask`

The default RPS CPU mask used on newly created network devices. An empty mask means RPS disabled by default.

CHAPTER 6. DEVICE DRIVERS

6.1. NEW DRIVERS

Network drivers

- Solarflare Siena network driver (**sfc-siena**), only in IBM Power Systems, Little Endian and AMD and Intel 64-bit architectures
- Nvidia sn2201 platform driver (**nvswn-sn2201**), only in AMD and Intel 64-bit architectures
- AMD SEV Guest Driver (**sev-guest**), only in AMD and Intel 64-bit architectures
- TDX Guest Driver (**tdx-guest**), only in AMD and Intel 64-bit architectures

Graphics drivers and miscellaneous drivers

- ACPI Video Driver (**video**), only in 64-bit ARM architecture
- DRM Buddy Allocator (**drm_buddy**), only in 64-bit ARM architecture and IBM Power Systems, Little Endian
- DRM display adapter helper (**drm_display_helper**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures
- Intel® GVT-g for KVM (**kvmgt**), only in AMD and Intel 64-bit architectures
- HP® iLO/iLO2 management processor (**hpilo**), only in 64-bit ARM architecture
- HPE watchdog driver (**hpwdt**), only in 64-bit ARM architecture
- AMD HSMP Platform Interface Driver (**amd_hsmpt**), only in AMD and Intel 64-bit architectures

6.2. UPDATED DRIVERS

Network drivers

- Intel® 10 Gigabit PCI Express Network Driver (**ixgbe**) has been updated to version 4.18.0-477 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Intel® 10 Gigabit Virtual Function Network Driver (**ixgbev**) has been updated to version 4.18.0-477 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Intel® 2.5G Ethernet Linux Driver (**igc**) has been updated to version 4.18.0-477 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Intel® Ethernet Adaptive Virtual Function Network Driver (**iaevf**) has been updated to version 4.18.0-477 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Intel® Ethernet Connection XL710 Network Driver (**i40e**) has been updated to version 4.18.0-477 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).

- Intel® Ethernet Switch Host Interface Driver (**fm10k**) has been updated to version 4.18.0-477 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Intel® Gigabit Ethernet Network Driver (**igb**) has been updated to version 4.18.0-477. (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Intel® Gigabit Virtual Function Network Driver (**igbvf**) has been updated to version 4.18.0-477 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Intel® PRO/1000 Network Driver (**e1000e**) has been updated to version 4.18.0-477 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Mellanox 5th generation network adapters (ConnectX series) core driver (**mlx5_core**) has been updated to version 4.18.0-477.
- The Netronome Flow Processor (NFP) driver (**nfp**) has been updated to version 4.18.0-477.

Storage drivers

- Driver for Microchip Smart Family Controller version (**smartpqi**) has been updated to version 2.1.20-035 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Emulex LightPulse Fibre Channel SCSI driver (**lpfc**) has been updated to version 14.0.0.18 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- LSI MPT Fusion SAS 3.0 Device Driver (**mpt3sas**) has been updated to version 43.100.00.00 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- MPI3 Storage Controller Device Driver (**mpi3mr**) has been updated to version 8.2.0.3.0 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- QLogic Fibre Channel HBA Driver(**qla2xxx**) has been updated to version 10.02.07.900-k (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- SCSI debug adapter driver (**scsi_debug**) has been updated to version 0191.

CHAPTER 7. AVAILABLE BPF FEATURES

This chapter provides the complete list of **Berkeley Packet Filter (BPF)** features available in the kernel of this minor version of Red Hat Enterprise Linux 8. The tables include the lists of:

- [System configuration and other options](#)
- [Available program types and supported helpers](#)
- [Available map types](#)

This chapter contains automatically generated output of the **bpftool feature** command.

Table 7.1. System configuration and other options

Option	Value
unprivileged_bpf_disabled	1 (bpf() syscall restricted to privileged users, without recovery)
JIT compiler	1 (enabled)
JIT compiler hardening	1 (enabled for unprivileged users)
JIT compiler kallsyms exports	1 (enabled for root)
Memory limit for JIT for unprivileged users	264241152
CONFIG_BPF	y
CONFIG_BPF_SYSCALL	y
CONFIG_HAVE_EBPF_JIT	y
CONFIG_BPF_JIT	y
CONFIG_BPF_JIT_ALWAYS_ON	y
CONFIG_DEBUG_INFO_BTF	y
CONFIG_DEBUG_INFO_BTF_MODULES	n
CONFIG_CGROUPS	y
CONFIG_CGROUP_BPF	y
CONFIG_CGROUP_NET_CLASSID	y
CONFIG_SOCK_CGROUP_DATA	y

Option	Value
CONFIG_BPF_EVENTS	y
CONFIG_KPROBE_EVENTS	y
CONFIG_UPROBE_EVENTS	y
CONFIG_TRACING	y
CONFIG_FTRACE_SYSCALLS	y
CONFIG_FUNCTION_ERROR_INJECTION	y
CONFIG_BPF_KPROBE_OVERRIDE	y
CONFIG_NET	y
CONFIG_XDP_SOCKETS	y
CONFIG_LWTUNNEL_BPF	y
CONFIG_NET_ACT_BPF	m
CONFIG_NET_CLS_BPF	m
CONFIG_NET_CLS_ACT	y
CONFIG_NET_SCH_INGRESS	m
CONFIG_XFRM	y
CONFIG_IP_ROUTE_CLASSID	y
CONFIG_IPV6_SEG6_BPF	n
CONFIG_BPF_LIRC_MODE2	n
CONFIG_BPF_STREAM_PARSER	y
CONFIG_NETFILTER_XT_MATCH_BPF	m
CONFIG_BPFILTER	n
CONFIG_BPFILTER_UMH	n

Option	Value
CONFIG_TEST_BPF	m
CONFIG_HZ	1000
bpf() syscall	available
Large program size limit	available

Table 7.2. Available program types and supported helpers

Program type	Available helpers
socket_filter	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_skb_load_bytes, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_skb_load_bytes_relative, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
kprobe	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_perf_event_read_value, bpf_override_return, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
sched_cls	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_l3_csum_replace, bpf_l4_csum_replace, bpf_tail_call, bpf_clone_redirect, bpf_get_cgroup_classid, bpf_skb_vlan_push, bpf_skb_vlan_pop, bpf_skb_get_tunnel_key, bpf_skb_set_tunnel_key, bpf_redirect, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_get_tunnel_opt, bpf_skb_set_tunnel_opt, bpf_skb_change_proto, bpf_skb_change_type, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_change_tail, bpf_skb_pull_data, bpf_csum_update, bpf_set_hash_invalid, bpf_get_numa_node_id, bpf_skb_change_head, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_set_hash, bpf_skb_adjust_room, bpf_skb_get_xfrm_state, bpf_skb_load_bytes_relative, bpf_fib_lookup, bpf_skb_cgroup_id, bpf_skb_ancestor_cgroup_id, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_fullsock, bpf_tcp_sock, bpf_skb_ecn_set_ce, bpf_get_listener_sock, bpf_skc_lookup_tcp, bpf_tcp_check_syncookie, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_tcp_gen_syncookie, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_sk_assign, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_csum_level, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_skb_cgroup_classid, bpf_redirect_neigh, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_redirect_peer, bpf_ktime_get_coarse_ns, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf
sched_act	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_l3_csum_replace, bpf_l4_csum_replace, bpf_tail_call, bpf_clone_redirect, bpf_get_cgroup_classid, bpf_skb_vlan_push, bpf_skb_vlan_pop, bpf_skb_get_tunnel_key, bpf_skb_set_tunnel_key, bpf_redirect, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_get_tunnel_opt, bpf_skb_set_tunnel_opt, bpf_skb_change_proto, bpf_skb_change_type, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_change_tail, bpf_skb_pull_data, bpf_csum_update, bpf_set_hash_invalid, bpf_get_numa_node_id, bpf_skb_change_head, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_set_hash, bpf_skb_adjust_room, bpf_skb_get_xfrm_state, bpf_skb_load_bytes_relative, bpf_fib_lookup, bpf_skb_cgroup_id, bpf_skb_ancestor_cgroup_id, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_fullsock, bpf_tcp_sock, bpf_skb_ecn_set_ce, bpf_get_listener_sock, bpf_skc_lookup_tcp, bpf_tcp_check_syncookie, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_tcp_gen_syncookie, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_sk_assign, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_csum_level, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_skb_cgroup_classid, bpf_redirect_neigh, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_redirect_peer, bpf_ktime_get_coarse_ns, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
tracepoint	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_perf_event_read_value, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
xdp	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_redirect, bpf_perf_event_output, bpf_csum_diff, bpf_get_current_task, bpf_get_numa_node_id, bpf_xdp_adjust_head, bpf_redirect_map, bpf_xdp_adjust_meta, bpf_xdp_adjust_tail, bpf_fib_lookup, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_lookup_tcp, bpf_tcp_check_syncookie, bpf_tcp_gen_syncookie, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf
perf_event	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_perf_event_read_value, bpf_perf_prog_read_value, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_read_branch_records, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
cgroup_skb	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_skb_load_bytes, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_skb_load_bytes_relative, bpf_skb_cgroup_id, bpf_get_local_storage, bpf_skb_ancestor_cgroup_id, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_fullsock, bpf_tcp_sock, bpf_skb_ecn_set_ce, bpf_get_listener_sock, bpf_skc_lookup_tcp, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_sk_cgroup_id, bpf_skb_ancestor_cgroup_id, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
cgroup_sock	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_get_cgroup_classid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_storage_get, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_netns_cookie, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
lwt_in	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_cgroup_classid, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_pull_data, bpf_get_numa_node_id, bpf_lwt_push_encap, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
lwt_out	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_cgroup_classid, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_pull_data, bpf_get_numa_node_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
lwt_xmit	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_l3_csum_replace, bpf_l4_csum_replace, bpf_tail_call, bpf_clone_redirect, bpf_get_cgroup_classid, bpf_skb_get_tunnel_key, bpf_skb_set_tunnel_key, bpf_redirect, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_get_tunnel_opt, bpf_skb_set_tunnel_opt, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_change_tail, bpf_skb_pull_data, bpf_csum_update, bpf_set_hash_invalid, bpf_get_numa_node_id, bpf_skb_change_head, bpf_lwt_push_encap, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_csum_level, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
sock_ops	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_setsockopt, bpf_sock_map_update, bpf_getsockopt, bpf_sock_ops_cb_flags_set, bpf_sock_hash_update, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_tcp_sock, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_load_hdr_opt, bpf_store_hdr_opt, bpf_reserve_hdr_opt, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
sk_skb	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_tail_call, bpf_perf_event_output, bpf_skb_load_bytes, bpf_get_current_task, bpf_skb_change_tail, bpf_skb_pull_data, bpf_get_numa_node_id, bpf_skb_change_head, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_skb_adjust_room, bpf_sk_redirect_map, bpf_sk_redirect_hash, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_skc_lookup_tcp, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
cgroup_device	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_uid_gid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
sk_msg	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_cgroup_classid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_msg_redirect_map, bpf_msg_apply_bytes, bpf_msg_cork_bytes, bpf_msg_pull_data, bpf_msg_redirect_hash, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_msg_push_data, bpf_msg_pop_data, bpf_spin_lock, bpf_spin_unlock, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
raw_tracepoint	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_perf_event_read_value, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
cgroup_sock_addr	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_get_cgroup_classid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_setsockopt, bpf_getsockopt, bpf_bind, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_skc_lookup_tcp, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_netns_cookie, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
lwt_seg6local	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_cgroup_classid, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_pull_data, bpf_get_numa_node_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
lirc_mode2	not supported

Program type	Available helpers
sk_reuseport	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_skb_load_bytes, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_skb_load_bytes_relative, bpf_sk_select_reuseport, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
flow_dissector	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_skb_load_bytes, bpf_get_current_task, bpf_get_numa_node_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
cgroup_sysctl	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_uid_gid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sysctl_get_name, bpf_sysctl_get_current_value, bpf_sysctl_get_new_value, bpf_sysctl_set_new_value, bpf_strtol, bpf_strtoul, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
raw_tracepoint_wri table	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_perf_event_read_value, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
cgroup_socketopt	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_uid_gid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_tcp_sock, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
tracing	not supported

Program type	Available helpers
struct_ops	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_l3_csum_replace, bpf_l4_csum_replace, bpf_tail_call, bpf_clone_redirect, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_get_cgroup_classid, bpf_skb_vlan_push, bpf_skb_vlan_pop, bpf_skb_get_tunnel_key, bpf_skb_set_tunnel_key, bpf_perf_event_read, bpf_redirect, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_get_stackid, bpf_csum_diff, bpf_skb_get_tunnel_opt, bpf_skb_set_tunnel_opt, bpf_skb_change_proto, bpf_skb_change_type, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_skb_change_tail, bpf_skb_pull_data, bpf_csum_update, bpf_set_hash_invalid, bpf_get_numa_node_id, bpf_skb_change_head, bpf_xdp_adjust_head, bpf_probe_read_str, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_set_hash, bpf_setsockopt, bpf_skb_adjust_room, bpf_redirect_map, bpf_sk_redirect_map, bpf_sock_map_update, bpf_xdp_adjust_meta, bpf_perf_event_read_value, bpf_perf_prog_read_value, bpf_getsockopt, bpf_override_return, bpf_sock_ops_cb_flags_set, bpf_msg_redirect_map, bpf_msg_apply_bytes, bpf_msg_cork_bytes, bpf_msg_pull_data, bpf_bind, bpf_xdp_adjust_tail, bpf_skb_get_xfrm_state, bpf_get_stack, bpf_skb_load_bytes_relative, bpf_fib_lookup, bpf_sock_hash_update, bpf_msg_redirect_hash, bpf_sk_redirect_hash, bpf_lwt_push_encap, bpf_lwt_seg6_store_bytes, bpf_lwt_seg6_adjust_srh, bpf_lwt_seg6_action, bpf_rc_repeat, bpf_rc_keydown, bpf_skb_cgroup_id, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_sk_select_reuseport, bpf_skb_ancestor_cgroup_id, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_msg_push_data, bpf_msg_pop_data, bpf_rc_pointer_rel, bpf_spin_lock, bpf_spin_unlock, bpf_sk_fullsock, bpf_tcp_sock, bpf_skb_ecn_set_ce, bpf_get_listener_sock, bpf_skc_lookup_tcp, bpf_tcp_check_syncookie, bpf_sysctl_get_name, bpf_sysctl_get_current_value, bpf_sysctl_get_new_value, bpf_sysctl_set_new_value, bpf_strtol, bpf_strtoul, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_send_signal, bpf_tcp_gen_syncookie, bpf_skb_output, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_tcp_send_ack, bpf_send_signal_thread, bpf_jiffies64, bpf_read_branch_records, bpf_get_ns_current_pid_tgid, bpf_xdp_output, bpf_get_netns_cookie, bpf_get_current_ancestor_cgroup_id, bpf_sk_assign, bpf_ktime_get_boot_ns, bpf_seq_printf, bpf_seq_write, bpf_sk_cgroup_id, bpf_sk_ancestor_cgroup_id, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_csum_level, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_get_task_stack, bpf_load_hdr_opt, bpf_store_hdr_opt, bpf_reserve_hdr_opt, bpf_inode_storage_get, bpf_inode_storage_delete, bpf_d_path, bpf_copy_from_user, bpf_snprintf_btf, bpf_seq_printf_btf, bpf_skb_cgroup_classid, bpf_redirect_neigh, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_redirect_peer, bpf_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_bprm_opts_set, bpf_ktime_get_coarse_ns, bpf_ima_inode_hash, bpf_sock_from_file, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf, bpf_sys_bpf, bpf_btf_find_by_name_kind, bpf_sys_close
ext	not supported
lsm	not supported

Program type	Available helpers
sk_lookup	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_sk_assign, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Table 7.3. Available map types

Map type	Available
hash	yes
array	yes
prog_array	yes
perf_event_array	yes
percpu_hash	yes
percpu_array	yes
stack_trace	yes
cgroup_array	yes
lru_hash	yes
lru_percpu_hash	yes
lpm_trie	yes
array_of_maps	yes
hash_of_maps	yes
devmap	yes
sockmap	yes

Map type	Available
cpumap	yes
xskmap	yes
sockhash	yes
cgroup_storage	yes
reuseport_sockarray	yes
percpu_cgroup_storage	yes
queue	yes
stack	yes
sk_storage	yes
devmap_hash	yes
struct_ops	no
ringbuf	yes
inode_storage	yes
task_storage	no

CHAPTER 8. BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 8.8 that have a significant impact on users.

8.1. INSTALLER AND IMAGE CREATION

Installer now lists all **PPC PreP Boot** or **BIOS Boot** partitions during custom partitioning

Previously, when adding multiple **PPC PreP Boot** or **BIOS Boot** partitions during custom partitioning, the Custom Partitioning screen displayed only one partition of a related type. As a consequence, the Custom Partitioning screen did not reflect the real state of the intended partitioning layout, making the partitioning process difficult and non-transparent.

With this update, the Custom Partitioning screen correctly displays all **PPC PreP Boot** or **BIOS Boot** partitions in the partitions list. As a result, users can now better understand and manage the intended partitioning layout.

[Bugzilla:1913035](#)

The installer now adds configuration options correctly into the yum repo files

Previously, the installer did not add configuration options correctly into yum repo files while including and excluding packages from additional installation repositories. With this update, yum repo files are created correctly. As a result, using the `--excludepkgs=` or `--includepkgs=` options in the `repo` kickstart command now excludes or includes the specified packages during installation as expected.

[Bugzilla:2014103](#)

Using the `filename DHCP` option no longer blocks downloading the `kickstart` file for installation

Previously, when building a path for getting the kickstart file from an NFS server, the installer did not consider the `filename DHCP` option. As a consequence, the installer did not download the kickstart file and was blocking the installation process. With this update, the `filename DHCP` option correctly constructs a path to the kickstart file. As a result, the kickstart file is downloaded properly, and the installation process starts correctly.

[Bugzilla:1991516](#)

The installer now creates a new GPT disk layout while custom partitioning

Previously, the installer did not change the disk layout to GPT when `inst.gpt` was specified on the kernel command line, and the user removed all partitions from a disk with the MBR disk layout on the custom partitioning spoke. As a consequence, the MBR disk layout remained on the disk.

With this update, the installer creates a new GPT disk layout on the disk if `inst.gpt` is specified on the kernel command line, and all partitions are removed from a disk on the custom partitioning spoke.

[Bugzilla:2094977](#)

The `--size` parameter of the `composer-cli compose start` command now treats its values as MiB

Previously, when using the `composer-cli compose start --size size_value blueprint_name image_type` command, the `composer-cli` tool treated the `--size` parameter values as byte units. This update fixes the issue, and the `--size` parameter values are now correctly used in the MiB format.

[Bugzilla:2033192](#)

8.2. SOFTWARE MANAGEMENT

RPM no longer hangs during a transaction involving the **fapolicyd** service restart

Previously, if you tried to update a package that caused the **fapolicyd** service to be restarted, for example, **systemd**, the RPM transaction stopped responding because the **fapolicyd** plug-in failed to communicate with the **fapolicyd** daemon.

With this update, the **fapolicyd** plug-in now correctly communicates with the **fapolicyd** daemon. As a result, RPM no longer hangs during a transaction which involves the **fapolicyd** service restart.

[Bugzilla:2110787](#)

Security YUM upgrade is now possible for packages that change their architecture through the upgrade

Patch for [BZ#2088149](#) introduced with [RHBA-2022:7711](#) caused a regression where YUM upgrade using security filters skipped packages that changed their architecture from or to **noarch** through the upgrade. Consequently, the missing security upgrades for these packages could leave the system in a vulnerable state.

With this update, the issue has been fixed, and security YUM upgrade no longer skips packages that change architecture from or to **noarch**.

[Bugzilla:2124483](#)

Reverting a YUM upgrade transaction is now possible for a package group or environment

Previously, the **yum history rollback** command failed when attempting to revert an upgrade transaction for a package group or an environment.

With this update, the issue has been fixed, and you can now revert the YUM upgrade transaction for a package group or environment.

[Bugzilla:2016070](#)

8.3. SHELLS AND COMMAND-LINE TOOLS

wsmancli handles HTTP 401 Unauthorized statuses correctly

The **wsmancli** utility for managing systems using Web Services Management protocol now handles authentication to better conform to RFC 2616.

Previously, when connecting to a service that requires authentication, the **wsmancli** command returned the error message **Authentication failed, please retry** immediately after receiving an HTTP 401 Unauthorized response, for example, because of incomplete credentials. To proceed, **wsmancli** prompted you to provide both the username and the password, even in situations where you had already provided a part of your credentials.

With this update, **wsmancli** requires only credentials that were not previously provided. As a result, the first authentication attempt does not display any error message. An error message is displayed only after you provide the complete credentials and authentication fails.

[Bugzilla:2105316](#)

The `translator.sty` LaTeX style document has been added

Previously, the `translator.sty` LaTeX style document, which is necessary for certain tools that depend on `texlive-beamer`, was missing. As a consequence, these tools failed with a **LaTeX Error: File `translator.sty` not found.** error. This update adds the missing `texlive-translator` package that contains the `translator.sty` LaTeX style document. As a result, tools that depend on `texlive-beamer` work correctly.

[Bugzilla:2150727](#)

ReaR handles excluded DASDs on the IBM Z architecture correctly

Previously on the IBM Z architecture, ReaR reformatted all connected Direct Access Storage Devices (DASD) during the recovery process, including those DASDs that users excluded from the saved layout and did not intend to restore their content. As a consequence, if you excluded some DASDs from the saved layout, their data were lost during system recovery. With this update, ReaR no longer formats excluded DASDs during system recovery, including the device from which the ReaR rescue system was booted (using the zIPL bootloader). You are also prompted to confirm the DASD formatting script before ReaR reformats DASDs. This ensures that the data on excluded DASDs survive a system recovery.

[Bugzilla:2172605](#)

ReaR no longer fails to restore non-LVM XFS filesystems

Previously, when you used ReaR to restore a non-LVM XFS filesystems with certain settings and disk mapping, ReaR created the file system with the default settings instead of the specified settings.

For example, if you had a file system with the `sunit` and `swidth` parameters set to non-zero values and you restored the file system using ReaR with disk mapping, the file system would be created with default `sunit` and `swidth` parameters ignoring the specified values.

As a consequence, ReaR failed during mounting the filesystem with specific XFS options. With this update, ReaR correctly restores the file system with the specified settings.

[Bugzilla:2131946](#)

8.4. INFRASTRUCTURE SERVICES

`rsync` no longer fails while using regular expressions for extended attributes

Previously, the `rsync` utility for transferring and synchronizing files was not able to handle extended attributes in RHEL 8 correctly. For example, if you passed the `--delete` option together with the `--filter 'x string.*'` option for extended attributes to the `rsync` command, and a file on your system satisfied the regular expression, this resulted in an error message stating protocol incompatibilities. With this update, the `rsync` utility handles extended attributes correctly and you can use regular expressions for these attributes.

[Bugzilla:2139118](#)

8.5. SECURITY

Scans and remediations correctly ignore SCAP Audit rules Audit key

Previously, Audit watch rules that were defined without an Audit key (`-k` or `-F` key) encountered the following problems:

- The rules were marked as non-compliant even if other parts of the rule were correct.
- Bash remediation fixed the path and permissions of the watch rule, but it did not add the Audit key correctly.
- Remediation sometimes did not fix the missing key, returning an **error** instead of a **fixed** value.

This affected the following rules:

- **audit_rules_login_events**
- **audit_rules_login_events_faillock**
- **audit_rules_login_events_lastlog**
- **audit_rules_login_events_tallylog**
- **audit_rules_usergroup_modification**
- **audit_rules_usergroup_modification_group**
- **audit_rules_usergroup_modification_gshadow**
- **audit_rules_usergroup_modification_opasswd**
- **audit_rules_usergroup_modification_passwd**
- **audit_rules_usergroup_modification_shadow**
- **audit_rules_time_watch_localtime**
- **audit_rules_mac_modification**
- **audit_rules_networkconfig_modification**
- **audit_rules_sysadmin_actions**
- **audit_rules_session_events**
- **audit_rules_sudoers**
- **audit_rules_sudoers_d**

With this update, the Audit key has been removed from checks and from Bash and Ansible remediations. As a result, inconsistencies caused by the key field during checking and remediating no longer occur, and auditors can choose these keys arbitrarily to make searching Audit logs easier.

[Bugzilla:2119356](#)

crypto-policies no longer creates unnecessary symlink

During system installation, the **crypto-policies** scriptlet creates symlinks from the **/usr/share/crypto-policies/DEFAULT** file or **/usr/share/crypto-policies/FIPS** in FIPS mode and saves them in the **/etc/crypto-policies/back-ends** directory. Previously, **crypto-policies** incorrectly included directories, and created a **/etc/crypto-policies/back-ends.config** symlink that pointed to the **/usr/share/crypto-policies/DEFAULT** or **/usr/share/crypto-policies/FIPS** directories. With this update, **crypto-policies** does not create symlinks from directories, and therefore does not create this unnecessary symlink.

[Bugzilla:1921646](#)

crypto-policies now disable NSEC3DSA for BIND

Previously, the system-wide cryptographic policies did not control the **NSEC3DSA** algorithm in the BIND configuration. Consequently, **NSEC3DSA**, which does not meet current security requirements, was not disabled on DNS servers. With this update, all cryptographic policies disable **NSEC3DSA** in the BIND configuration by default.

[Bugzilla:2071981](#)

Libreswan no longer rejects SHA-1 signature verification in the FUTURE and FIPS cryptographic policies

Previously, from update to 4.9, Libreswan rejected SHA-1 signature verification in the **FUTURE** and **FIPS** cryptographic policies, and peer authentication failed when **authby=rsasig** or **authby=rsa-sha1** connection options were used. This update reverts this behavior by relaxing how Libreswan handles the **crypto-policies** settings. As a consequence, you can now use the **authby=rsasig** and **authby=rsa-sha1** connection options using SHA-1 signature verification.

[Bugzilla:2176248](#)

crontab bash scripts no longer execute in incorrect context

Previously, a bug fix published in erratum [RHBA-2022:7691](#) used too general transition rule. Consequently, a bash script executed from the **crontab** file was executed in the **rpm_script_t** context instead of the **system_cronjob_t** context. With this update, bash scripts are now executed in the correct context.

[Bugzilla:2154242](#)

selinux-policy supports service execution in SAP Host Agent

Previously, the SELinux policy did not support the **insights-client** service interacting with SAP Host Agent and other services. As a consequence, some commands did not work correctly when started from Red Hat Insights. With this update, the SELinux policy supports SAP service execution. As a result, SAP services started from Insights run successfully.

[Bugzilla:2134125](#)

selinux-policy now allows pmcd to execute its private memfd: objects

Previously, the SELinux policy did not allow the **pmcd** process from the Performance Co-Pilot (PCP) framework to execute its private memory file-system objects (**memfd:**). Consequently, SELinux denied the Performance Metric Domain Agent (PMDA) BPF Compiler Collection (BCC) service to execute **memfd:** objects. In this update, the SELinux policy contains new rules for **pmcd**. As a result, **pmcd** can now execute **memfd:** objects with SELinux in enforcing mode.

[Bugzilla:2090711](#)

SELinux policy allows sysadm_r to use subscription-manager

Previously, users in the **sysadm_r** SELinux role were not allowed to execute some subcommands of the **subscription-manager** utility. Consequently, the subcommands failed to read the memory device. This update adds a new rule to the SELinux policy that allows the **sysadm_t** type to read **/dev/mem**. As a consequence, the **subscription-manager** subcommands do not fail.

[Bugzilla:2101341](#)

samba-dcerpcd process now works correctly with nscd

Previously, the **samba-dcerpcd** process could not communicate with the **nscd** process because of the SELinux policy. Consequently, the **samba-dcerpcd** service did not work properly when the **nscd** service was enabled. With this update, the SELinux policy has been updated with new rules for **samba-dcerpcd**.

[Bugzilla:2121709](#)

vlock now works properly for confined users

Previously, the confined user could not use **vlock** due to SELinux policy. Consequently, the **vlock** command did not work properly for confined users. With this update, the SELinux policy has been updated with new rules for confined users.

[Bugzilla:2122838](#)

Confined users now can log in without a reported denial

Previously, SELinux policy did not allow all permissions needed to log in a SELinux confined user using GUI. Consequently, AVC denials were audited and some services like **dbus** or **pulseaudio** did not work properly. With this update, the SELinux policy has been updated with new rules for confined users.

[Bugzilla:2124388](#)

insights-client now has additional permissions in the SELinux policy

The updated **insights-client** service requires additional permissions, which were not included in the previous versions of the **selinux-policy** packages. As a consequence, certain components of **insights-client** did not work correctly with SELinux in enforcing mode, and the system reported access vector cache (AVC) error messages. This update adds the missing permissions to the SELinux policy. As a result, **insights-client** now works correctly without reporting AVC errors.

[Bugzilla:2125008](#)

The SELinux policy allows smb access to user shares

Previously, the **samba-dcerpcd** process was separated from the **smb** service, but did not have access to user shares. As a consequence, **smb** clients could not access files on user **smb** shares. This update adds rules to the SELinux policy for managing user home content for the **samba-dcerpcd** binary when the **samba_enable_home_dirs** boolean is enabled. As a result, **samba-dcerpcd** can access user shares when **samba_enable_home_dirs** is on.

[Bugzilla:2143696](#)

The SELinux policy now allows confined administrators to access ipmi devices when IPMItool runs

Previously, the SELinux policy did not allow confined administrators to read and write **ipmi** devices when the IPMItool utility is run. As a consequence, when a confined administrator ran **ipmitool**, it failed. This update adds allow rules to **selinux-policy** for administrators assigned to the **sysadm_r** SELinux role. As a result, if a confined administrator runs **ipmitool** it works correctly.

[Bugzilla:2148561](#)

SCAP Security Guide rule file_permissions_sshd_private_key is aligned with STIG configuration RHEL-08-010490

Previously, the implementation of rule **file_permissions_sshd_private_key** allowed private SSH keys to be readable by the **ssh_keys** group with mode **0644**, while DISA STIG version RHEL-08-010490

required private SSH keys to have mode **0600**. As a consequence, evaluation with DISA's automated STIG benchmark failed for configuration RHEL-08-010490.

For this update, we worked with DISA to align the expected permissions for private SSH keys, and now private keys are expected to have mode **0644** or less permissive. As a result, the rule **file_permissions_sshd_private_key** and configuration RHEL-08-010490 are now aligned.

[Bugzilla:2115343](#)

The **sudo_require_reauthentication** SCAP Security Guide rule accepts correct spacing in **sudoers**

Previously, a bug in the checking of the **xccdf_org.ssgproject.content_rule_sudo_require_reauthentication** rule caused it to require specific spacing between the **timestamp_timeout** key and its value in the **/etc/sudoers** file and the **/etc/sudoers.d** directory. Consequently, valid and compliant syntax caused the rule to fail incorrectly. With this update, the check for **xccdf_org.ssgproject.content_rule_sudo_require_reauthentication** has been updated to accept blank spaces around the equal sign. As a result, the rule accepts correct and compliant definitions of **timestamp_timeout** with any of the following spacing formats:

- **Defaults timestamp_timeout = 5**
- **Defaults timestamp_timeout= 5**
- **Defaults timestamp_timeout =5**
- **Defaults timestamp_timeout=5**

[Bugzilla:2152208](#)

Old Kerberos rules changed to **notapplicable** in new versions of RHEL

Previously, some Kerberos-related rules failed while scanning against the DISA STIG profile on RHEL 8.8 and later systems in FIPS mode, even though the system should have been compliant. This was caused by the following rules:

- **xccdf_org.ssgproject.content_rule_package_krb5-server_removed**
- **xccdf_org.ssgproject.content_rule_package_krb5-workstation_removed**
- **xccdf_org.ssgproject.content_rule_kerberos_disable_no_keytab**

This update makes these rules not applicable for RHEL versions 8.8 and later. As a result, the scan correctly returns the **notapplicable** result for these rules.

[Bugzilla:2099394](#)

scap-security-guide STIG profiles no longer require specific text in **/etc/audit/rules.d/11-loginuid.rules**

Previously, the SCAP rule **audit_immutable_login_uids** used in RHEL 8 profiles **stig** and **stig_gui** passed only if file **/etc/audit/rules.d/11-loginuid.rules** contained exact text. This is, however, not necessary to fulfill the STIG requirement (RHEL-08-030122). With this update, the new rule **audit_rules_immutable_login_uids** replaces **audit_immutable_login_uids** in RHEL 8 **stig** and **stig_gui** profiles. As a result, you can now specify the **--loginuid-immutable** parameter that fulfills the rule in any file with the **.rules** extension within the **/etc/audit/rules.d** directory or in the **/etc/audit/audit.rules** file, depending on usage of **auditctl** or **augen-rules**.

[Bugzilla:2151553](#)

Rules for CIS profiles in `scap-security-guide` are better aligned

Previously, some rules were incorrectly assigned to certain Center for Internet Security (CIS) profiles (`cis`, `cis_server_l1`, `cis_workstation_1`, and `cis_workstation_l2`). As a consequence, scanning according to some CIS profiles could skip rules from the CIS benchmark or check for unnecessary rules.

The following rules were assigned to incorrect profiles:

- Rules `kernel_module_udf_disabled`, `sudo_require_authentication` and `kernel_module_squashfs_disabled` were incorrectly placed in CIS Server Level 1 and CIS Workstation Level 1.
- Rules `package_libselinux_installed`, `grub2_enable_selinux`, `selinux_policytype`, `selinux_confinement_of_daemons`, `rsyslog_nolisten`, `service_systemd-journald_enabled` were missing from CIS Server Level 1 and CIS Workstation Level 1 profiles.
- Rules `package_setroubleshoot_removed` and `package_mcstrans_removed` were missing from the CIS Server Level 1 profile.

This update assigns the misaligned rules to the correct CIS profiles, but does not introduce new rules or entirely removes any rules. As a result, SCAP CIS profiles are better aligned with the original CIS benchmark.

[Bugzilla:2162803](#)

Clevis ignores commented devices in `crypttab`

Previously, Clevis tried to unlock commented devices in the `crypttab` file, causing the `clevis-luks-askpass` service to run even if the device was not valid. This caused unnecessary service runs and made it difficult to troubleshoot.

With this fix, Clevis ignores commented devices. Now, if an invalid device is commented, Clevis does not attempt to unlock it and `clevis-luks-askpass.service` finishes appropriately. This makes it easier to troubleshoot and reduces unnecessary service runs.

[Bugzilla:2159440](#)

Clevis no longer requests too much entropy from `pwmake`

Previously, the `pwmake` password generation utility displayed unwanted warnings when Clevis used `pwmake` to create passwords for storing data in `LUKS` metadata, which caused Clevis to use lower entropy. With this update, Clevis is limited to 256 entropy bits provided to `pwmake`, which eliminates an unwanted warning and uses the correct amount of entropy.

[Bugzilla:2159736](#)

`logrotate` no longer incorrectly signals Rsyslog in log rotation

Previously, the argument order was incorrectly set in the `logrotate` script, which caused a syntax error. This resulted in `logrotate` not correctly signaling Rsyslog during log rotation.

With this update, the order of the arguments in `logrotate` is fixed and `logrotate` signals Rsyslog correctly after log rotation even when the `POSIXLY_CORRECT` environment variable is set.

[Bugzilla:2070496](#)

Rsyslog no longer crashes due to a bug in `imklog`

Previously, Rsyslog could encounter a segmentation fault if the `imklog` module was enabled and a `free()` call using an invalid object was freed during use. With this update, the freed object is correctly deallocated at the correct place. As a result, the segmentation fault no longer occurs.

[Bugzilla:2157658](#)

USBGuard no longer causes a confusing warning

Previously, a race condition could happen in USBGuard when a parent process finished sooner than the first child process. As a consequence, `systemd` reported that a process was present with a wrongly identified parent PID (PPID). With this update, a parent process waits for the first child process to finish in working mode. As a result, `systemd` no longer reports such warnings.

[Bugzilla:2159409](#)

The `usbguard` service file did not define `OOMScore`

Previously, the `usbguard` service file did not define the `OOMScoreAdjust` option. Consequently, the process could be identified as a candidate for killing before unprivileged processes when the system resources are closed to running out. With this update, the new `OOMScoreAdjust` setting was introduced to the `usbguard.service` file, to disable OOM killing processes of the usbguard unit.

[Bugzilla:2159411](#)

USBGuard saves rules even if `RuleFile` is not defined

Previously, if the `RuleFile` configuration directive in USBGuard was set but `RuleFolder` was not, the rule set could not be changed. With this update, you can now change the rule set even if `RuleFolder` is set but `RuleFile` is not. As a result, you can modify the permanent policy in USBGuard to permanently save newly added rules.

[Bugzilla:2159413](#)

8.6. NETWORKING

`xdp-tools` rebased to version 1.2.10

The `xdp-tools` packages have been upgraded to upstream version 1.2.10, which provides a number of bug fixes over the previous version.

[Bugzilla:2160069](#)

`contrackd` functions properly even if `HashSize` and `HashLimit` are not set manually

Previously, the `contrackd` service did not set default values for the `HashSize` and `HashLimit` configuration variables. Consequently, `contrackd` could become unstable or stop functioning entirely if you did not specify those values. The problem has been fixed by making the configuration reader set the default values for `HashSize` and `HashLimit` before `contrackd` parses the configuration file. As a result, `contrackd` now functions correctly even if you do not specify the values.

[Bugzilla:2126736](#)

The `nm-cloud-setup` service no longer removes routes and manually-configured secondary IP addresses from interfaces

Based on the information received from the cloud environment, the `nm-cloud-setup` service configures

network interfaces. Previously, administrators had to disable **nm-cloud-setup** to manually configure routes and secondary IP addresses on interfaces to avoid that the service removes them. This update adds a flag to the **Reapply()** function to preserve externally added addresses and routes. As a result, administrators no longer need to disable the **nm-cloud-setup** service in the mentioned scenario.

[Bugzilla:2132754](#)

8.7. KERNEL

kpatch-patch works correctly on systems with an idle isolated CPU

Previously, when you attempted to install **kpatch-patch** CVE mitigation packages on systems with the kernel CPU isolation feature, the **kpatch-patch** RPMs did install, but failed to load their CVE mitigation kernel module. With this fix, the two features co-exist, and you can now successfully deploy **kpatch** CVE fixes when CPU isolation is in place.

[Bugzilla:2134931](#)

Enabling VMD works again

Previously, the operating system would fail to boot if Volume Management Device (VMD) was enabled. This update provides numerous bug fixes essential for VMD to work as expected.

[Bugzilla:2127028](#)

8.8. FILE SYSTEMS AND STORAGE

System works correctly without the soft lockup while starting a VDO volume

Due to fixing a Kernel Application Binary Interface (kABI) bug in the **pv_mmu_ops** structure, RHEL 8.7 systems with kernel version **4.18.0-425.10.1.el8_7**, that is RHEL-8.7.0.2-BaseOS, hung or encountered a kernel panic due to soft lockup while starting a Virtual Data Optimizer (VDO) volume.

With this update, the **kmod-kvdo** package was rebuilt any time a new kernel was available that is no longer kABI compatible with the current version of **kmod-kvdo**. As a result, the system works correctly while starting a VDO volume.

[Bugzilla:2119819](#)

VDO driver bug no longer causing device freezes through journal blocks

Previously, a bug in the VDO driver caused the system to mark some journal blocks as waiting for metadata updates. This problem was triggered when increasing the size of the VDO pool or the logical volume on top of it, or when using the **pvmove** and **lvchange** operations on LVM tools managed VDO devices. The bug was caused by incomplete resets that left some journal pages unavailable for use, and an incorrect notion of how many slots in the recovery journal were available to be filled. As a result, the device would freeze.

This issue has now been fixed with the latest version of the kernel modules for the virtual data optimizer **kmod-kvdo-6.2.8.1-87.el8**. Currently, all incomplete metadata blocks are saved in each section of the code in phases, while also updating in-memory data structures and resetting state on resume if needed. With this fix, users should no longer experience device freezes due to this issue.

[Bugzilla:2109047](#)

8.9. HIGH AVAILABILITY AND CLUSTERS

pcs no longer allows you to modify cluster properties that should not be changed

Previously, the **pcs** command line interface allowed you to modify cluster properties that should not be changed or for which change does not take effect. With this fix, **pcs** no longer allows you to modify these cluster properties: **cluster-infrastructure**, **cluster-name**, **dc-version**, **have-watchdog**, and **last-lrm-refresh**.

[Bugzilla:2112263](#)

pcs now displays cluster properties that are not explicitly configured

Previously, a **pcs** command to display the value of a specific cluster property did not list values that are not explicitly configured in the CIB. With this fix, if a cluster property is not set **pcs** displays the default value for the property.

[Bugzilla:2112267](#)

Cluster resources that call `crm_mon` now stop cleanly at shutdown

Previously, the **crm_mon** utility returned a nonzero exit status while Pacemaker was in the process of shutting down. Resource agents that called **crm_mon** in their monitor action, such as **ocf:heartbeat:pqsq**, could incorrectly return a failure at cluster shutdown. With this fix, **crm_mon** returns success even if the cluster is in the process of shutting down. Resources that call **crm_mon** now stop cleanly at cluster shutdown.

[Bugzilla:2133497](#)

OCF resource agent metadata actions can now call `crm_node` without causing unexpected fencing

As of RHEL 8.5, OCF resource agent metadata actions blocked the controller and **crm_node** queries performed controller requests. As a result, if an agent's metadata action called **crm_node**, it blocked the controller for 30 seconds until the action timed out. This could cause other actions to fail and the node to be fenced.

With this fix, the controller now performs metadata actions asynchronously. An OCF resource agent metadata action can now call **crm_node** without issue.

[Bugzilla:2121852](#)

Enabling a single resource and monitoring operation no longer enables monitoring operations for all resources in a resource group

Previously, after unmanaging all resources and monitoring operations in a resource group, managing one of the resources in that group along with its monitoring operation re-enabled the monitoring operations for all resources in the resource group. This could trigger unexpected cluster behavior.

With this fix, managing a resource and re-enabling its monitoring operation re-enables the monitoring operation for that resource only and not for the other resources in a resource group.

[Bugzilla:1918527](#)

Pacemaker now rechecks resource assignments immediately when resource order changes

As of RHEL 8.7, Pacemaker did not recheck resource assignments when the order of resources in the CIB changed with no changes to the resource definition. If configuration reordering would cause resources to move, that would not take place until the next natural transition, up to the value of **cluster-recheck-interval-property**. This could cause issues if resource stickiness is not configured for a resource.

With this change, Pacemaker rechecks resource assignments when the order of the resources in the CIB changes, as it did for earlier Pacemaker releases. The cluster now responds immediately to these changes, if needed.

[Bugzilla:2122806](#)

8.10. COMPILERS AND DEVELOPMENT TOOLS

You can install SciPy using pip on all architectures

Previously, the **openblas-devel** package did not contain a pkg-config file for the OpenBLAS library. As a consequence, in certain scenarios, it was impossible to determine the compiler and linker flags using the **pkgconf** utility while compiling with OpenBLAS. For example, this caused a failure of the **pip install scipy** command on the 64-bit IBM Z and IBM Power Systems, Little Endian architectures.

This update adds the **openblas.pc** file to the **openblas-devel** package on all supported architectures. As a result, you can install the SciPy library using the **pip** package installer.

[Bugzilla:2115722](#)

Functions in go no longer cause memory leak

Previously, the **EVP_PKEY_sign_raw** and **EVP_PKEY_verify_raw** functions did not call free to clean the memory. Consequently, the memory leaked and has not been recovered. With this updated, the **EVP_PKEY_sign_raw** and **EVP_PKEY_verify_raw** functions now call free and memory is not leaking.

[Bugzilla:2132767](#)

golang now supports 4096 bit keys in x509 FIPS mode

Previously, **golang** did not support the 4096 bit keys in x509 FIPS mode. Consequently, when the user used 4096 bit keys the program crashed. With this update, **golang** now supports 4096 bit keys in x509 FIPS mode.

[Bugzilla:2132694](#)

libffi can now probe for executable memory with SELinux enabled

By default, **libffi** does not probe for executable memory when SELinux is enabled. As a consequence, programs which use **libffi** closures and **fork()** without immediately executing some other processes terminate unexpectedly when SELinux is enabled. With this update, **libffi** looks for a **/etc/sysconfig/libffi-force-shared-memory-check-first** file and, if it exists, probes for executable memory regardless of if SELinux is enabled. As a result, programs using **libffi** can safely **fork()** without crashing with SELinux enabled.

[Bugzilla:2014228](#)

Implemented big endian support in OpenSSL bindings for golang

Previously, the **OpenSSL** bindings for **golang** did not have support for big-endian, leading to potential issues with the conversion of **BigInt** values. As a result, the crypto routines were unable to perform this conversion. To fix this issue, big-endian support was implemented in the **OpenSSL** bindings for **golang**. As a result, conversions from **BigInt** are now successful, and the tests pass as expected.

[Bugzilla:2132419](#)

8.11. IDENTITY MANAGEMENT

Authentication to external IdPs that require a client secret is now possible

Previously, SSSD did not properly pass client secrets to external identity providers (IdPs). Consequently, authentication failed against external IdPs that you previously configured with the **ipa idp-add --secret** command to require a client secret. With this update, SSSD passes the client secret to the IdP and users can authenticate.

Jira:RHELPLAN-148303

IdM now supports setting hostmasks for `sudo` rules using Ansible

Previously, the **ipa sudorule-add-host** command allowed setting a hostmask to be used by the `sudo` rule, but this option was not present in the **ansible-freeipa** package. With this update, you can now use the **ansible-freeipa hostmask** variable to define a list of hostmasks to which a particular `sudo` rule, defined in Identity Management (IdM), applies.

As a result, you can now automate setting host masks for IdM `sudo` rules with Ansible.

[Bugzilla:2127912](#)

The scheduled time of the changelog compaction now works correctly

Previously, when you configured a custom scheduled time for the changelog compaction, the server did not apply the new setting, and the changelog compaction could start during peak times. With this release, the server now correctly applies the custom time of the changelog compaction.

[Bugzilla:2130276](#)

IdM clients correctly retrieve information for trusted AD users when their names contain mixed case characters

Previously, if you attempted a user lookup or authentication of a user, and that trusted Active Directory (AD) user contained mixed case characters in their names and they were configured with overrides in IdM, an error was returned preventing users from accessing IdM resources.

With the release of [RHBA-2023:4525](#), a case-sensitive comparison is replaced with a case-insensitive comparison that ignores the case of a character. As a result, IdM clients can now lookup users of an AD trusted domain, even if their usernames contain mixed case characters and they are configured with overrides in IdM.

Jira:SSSD-6096

8.12. GRAPHICS INFRASTRUCTURE

Matrox G200e now works correctly with a VGA display

Previously, your display might have shown no graphical output if you used the following system configuration:

- The Matrox G200e GPU
- A display connected over the VGA controller

As a consequence, you could not use or install RHEL on this configuration.

With this release, the problem has been fixed. As a result, RHEL boots and shows graphical output as expected.

Bugzilla:2130159

8.13. THE WEB CONSOLE

The web console NBDE binding steps now work also on volume groups with a root file system

In RHEL 8.8.0, due to a bug in the code for determining whether or not the user was adding a Tang key to the root file system, the binding process in the web console crashed when there was no file system on the LUKS container at all. Because the web console displayed the error message **TypeError: Qe(...) is undefined** after you had clicked the **Trust key** button in the **Verify key** dialog, you had to perform all the required steps in the command-line interface in the described scenario.

With the release of the [RHBA-2023:3829](#) advisory, the web console correctly handles additions of Tang keys to root file systems. As a result, the web console finishes all binding steps required for the automated unlocking of LUKS-encrypted volumes using Network-Bound Disk Encryption (NBDE) in various scenarios.

Bugzilla:2212371

8.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The `nbde_client` System Role now correctly handles different names of `clevis-luks-askpass`

The `nbde_client` System Role has been updated to handle the systems on which the `clevis-luks-askpass systemd` unit has a different name. The role now correctly works with different names of `clevis-luks-askpass` on managed nodes, which requires unlocking also LUKS-encrypted volumes that mount late in the boot process.

Bugzilla:2126960

The `ha_cluster` System Role logs no longer display unencrypted passwords and secrets

The `ha_cluster` System Role accepts parameters that can be passwords or other secrets. Previously, some of the tasks would log their inputs and outputs. As a result, the role logs could contain unencrypted passwords and other secrets.

With this update, the tasks have been changed to use the Ansible `no_log: true` directive and the task output is no longer displayed in the role logs. The `ha_cluster` System Role logs no longer contain passwords and other secrets. While this update protects secure information, the role logs now provide less information that you can use when debugging your configuration.

Bugzilla:2127497

Clusters configured with `ha_cluster` System Role to use SBD and not start on boot now work correctly

Previously, if a user configured a cluster using the `ha_cluster` System Role to use SBD and not start on boot, then the SBD service was disabled and SBD did not start. With this fix, the SBD service is always enabled if a cluster is set to use SBD whether or not the cluster is configured to start on boot.

Bugzilla:2153081

Setting `stonith-watchdog-timeout` property with the `ha_cluster` System Role now works in a stopped cluster

Previously, when you set the **stonith-watchdog-timeout** property with the **ha_cluster** System Role in a stopped cluster, the property reverted to its previous value and the role failed. With this fix, configuring the **stonith-watchdog-timeout** property by using the **ha_cluster** System Role works properly.

[Bugzilla:2167941](#)

Enabling implicit files provider to fix **rhel-system-roles** SSSD configuration

A disabled SSSD implicit files provider caused the **rhel-system-roles** modules to create an invalid System Security Services Daemon (SSSD) configuration. This update unconditionally enables the files provider and as a result, the SSSD configuration created by **rhel-system-roles** now works as expected.

[Bugzilla:2153080](#)

Network traffic is now directed through the intended network interface when using **initscripts** with the **networking** RHEL System Role

Previously, when using the **initscripts** provider, the routing configuration for network connections did not specify the output device that the traffic should go through. Consequently, the kernel could use a different output device than the user intended. Now, if the network interface name is specified in the playbook for the connection, it is used as the output device in the route configuration file. This aligns the behavior with NetworkManager, which configures the output device in routes when activating profiles on devices. As a result, the users can ensure that the traffic is directed through the intended network interface.

[Bugzilla:2168733](#)

The **nbde_client_clevis** role no longer reports traceback to users

Previously, the **nbde_client_clevis** role sometimes failed in exception, causing a traceback and reporting sensitive data, such as the **encryption_password** field, back to the user. With this update, the role no longer reports sensitive data, only the appropriate error messages.

[Bugzilla:2162782](#)

8.15. VIRTUALIZATION

System time on nested VMs now works reliably

Previously, system time on nested virtual machines (VMs) in some cases desynchronised from the Level 0 and level 1 hosts. This also sometimes caused the nested VM to become unresponsive or terminate unexpectedly.

With this update, the time handling code in the KVM host kernel code has been fixed, which prevents the described errors from occurring.

[Bugzilla:2151854](#)

Network traffic performance in virtual machines is no longer reduced

Previously, RHEL virtual machines had, in some cases, decreased performance when handling high levels of network traffic. The underlying code has been fixed and the network traffic performance is not affected anymore.

[Bugzilla:2069047](#)

Virtual machines using **memfd** run as expected

Previously, virtual machines (VMs) running on the 64-bit IBM Z processor architecture that used **memfd** to back memory with hugepages failed to run. With this update, the problem has been fixed and VMs using **memfd** can now be defined on the 64-bit IBM Z processor architecture. As a result, you can now run VMs which use **memfd** to back the memory with hugepages.

[Bugzilla:2117149](#)

System time in VMs now synchronizes correctly with the host

Previously, the KVM module performed the real-time clock (RTC) synchronization less frequently than intended. As a consequence, the system time in VMs hosted on RHEL 8 sometimes did not correctly reflect the system time on the host. This update fixes the RTC scheduling in KVM, which prevents the described problem from occurring.

[Bugzilla:2135417](#)

CHAPTER 9. TECHNOLOGY PREVIEWS

This part provides a list of all Technology Previews available in Red Hat Enterprise Linux 8.8.

For information on Red Hat scope of support for Technology Preview features, see [Technology Preview Features Support Scope](#).

9.1. INFRASTRUCTURE SERVICES

Socket API for TuneD available as a Technology Preview

The socket API for controlling TuneD through Unix domain socket is now available as a Technology Preview. The socket API maps one-to-one with the D-Bus API and provides an alternative communication method for cases where D-Bus is not available. By using the socket API, you can control the TuneD daemon to optimize the performance, and change the values of various tuning parameters. The socket API is disabled by default, you can enable it in the **tuned-main.conf** file.

[Bugzilla:2113900](#)

9.2. NETWORKING

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.

[Bugzilla:1633143](#)

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 hardware offloading.

[Bugzilla:1889737](#)

Multi-protocol Label Switching for TC available as a Technology Preview

The Multi-protocol Label Switching (MPLS) is an in-kernel data-forwarding mechanism to route traffic flow across enterprise networks. In an MPLS network, the router that receives packets decides the further route of the packets based on the labels attached to the packet. With the usage of labels, the MPLS network has the ability to handle packets with particular characteristics. For example, you can add **tc filters** for managing packets received from specific ports or carrying specific types of traffic, in a consistent way.

After packets enter the enterprise network, MPLS routers perform multiple operations on the packets, such as **push** to add a label, **swap** to update a label, and **pop** to remove a label. MPLS allows defining actions locally based on one or multiple labels in RHEL. You can configure routers and set traffic control (**tc**) filters to take appropriate actions on the packets based on the MPLS label stack entry (**lse**) elements, such as **label**, **traffic class**, **bottom of stack**, and **time to live**.

For example, the following command adds a filter to the *enp0s1* network interface to match incoming packets having the first label *12323* and the second label *45832*. On matching packets, the following actions are taken:

- the first MPLS TTL is decremented (packet is dropped if TTL reaches 0)
- the first MPLS label is changed to *549386*
- the resulting packet is transmitted over *enp0s2*, with destination MAC address *00:00:5E:00:53:01* and source MAC address *00:00:5E:00:53:02*

```
# tc filter add dev enp0s1 ingress protocol mpls_uc flower mpls lse depth 1 label 12323 lse
depth 2 label 45832 \
action mpls dec_ttl pipe \
action mpls modify label 549386 pipe \
action pedit ex munge eth dst set 00:00:5E:00:53:01 pipe \
action pedit ex munge eth src set 00:00:5E:00:53:02 pipe \
action mirrored egress redirect dev enp0s2
```

Bugzilla:1814836, [Bugzilla:1856415](#)

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.

Bugzilla:1839311

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, a 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.

[Bugzilla:1906489](#)

KTLS available as a Technology Preview

RHEL provides Kernel Transport Layer Security (KTLS) 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 includes the interface for offloading TLS record encryption to Network Interface Controllers (NICs) that provides this functionality.

Bugzilla:1570255

9.3. KERNEL

Soft-RoCE available as a Technology Preview

Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) is a network protocol that implements RDMA over Ethernet. Soft-RoCE is the software implementation of RoCE which maintains two protocol versions, RoCE v1 and RoCE v2. The Soft-RoCE driver, **rdma_rxe**, is available as an

unsupported Technology Preview in RHEL 8.

Bugzilla:1605216

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 enables 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)** manual 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. 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:

- **AF_XDP**, a socket for connecting the **eXpress Data Path (XDP)** path to user space for applications that prioritize packet processing performance.

Bugzilla:1559616

The **kexec** fast reboot feature is available as a Technology Preview

The **kexec** fast reboot feature continues to be available as a Technology Preview. The **kexec** fast reboot significantly speeds the boot process as you can boot directly into the second kernel without passing through the Basic Input/Output System (BIOS) or firmware first. To use this feature:

1. Load the **kexec** kernel manually.
2. Reboot for changes to take effect.

Note that the **kexec** fast reboot capability is available with a limited scope of support on RHEL 9 and later releases.

[Bugzilla:1769727](#)

The Intel data streaming accelerator driver for kernel is available as a Technology Preview

The Intel data streaming accelerator driver (IDXD) for the kernel is currently available as a Technology Preview. It is an Intel CPU integrated accelerator and includes a shared work queue with process address space ID (pasid) submission and shared virtual memory (SVM).

Bugzilla:1837187

The **accel-config** package available as a Technology Preview

The **accel-config** package is now available on Intel **EM64T** and **AMD64** architectures as a Technology Preview. This package helps in controlling and configuring data-streaming accelerator (DSA) subsystem in the Linux Kernel. Also, it configures devices through **sysfs** (pseudo-file system), saves and loads the configuration in the **json** format.

Bugzilla:1843266

SGX available as a Technology Preview

Software Guard Extensions(SGX) is an Intel® technology for protecting software code and data from disclosure and modification. The RHEL kernel partially provides the SGX v1 and v1.5 functionality. Version 1 enables platforms using the **Flexible Launch Control** mechanism to use the SGX technology. Version 2 adds **Enclave Dynamic Memory Management**(EDMM). Notable features include:

- Modifying EPCM permissions of regular enclave pages that belong to an initialized enclave.
- Dynamic addition of regular enclave pages to an initialized enclave.
- Expanding an initialized enclave to accommodate more threads.
- Removing regular and TCS pages from an initialized enclave.

Bugzilla:1660337

9.4. FILE SYSTEMS AND STORAGE

File system DAX is now available for ext4 and XFS as a Technology Preview

In Red Hat Enterprise Linux 8, the 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 provides the capability of DAX must be created on the NVDIMM(s). Also, the file system must be mounted with the **dax** mount option. Then, a **mmap** of a file on the dax-mounted file system results in a direct mapping of storage into the application's address space.

Bugzilla: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 or other specialized use cases, such as squashed **kdump** initramfs. Its use is supported primarily 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](#).

For more information about OverlayFS, see the [Linux kernel documentation](#).

Bugzilla:1690207

Stratis is now available as a Technology Preview

Stratis is a new local storage manager, which provides managed file systems on top of pools of storage with additional features. It is provided as a Technology Preview.

With Stratis, you can perform the following storage tasks:

- 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. For more information, see the [Setting up Stratis file systems](#) documentation.

RHEL 8.5 updated Stratis to version 2.4.2. For more information, see the [Stratis 2.4.2 Release Notes](#).

Jira:RHELPLAN-1212

NVMe/TCP host 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** kernel module has been added as a Technology Preview. The use of NVMe/TCP as a host is manageable with tools provided by the **nvme-cli** package. The NVMe/TCP host Technology Preview is included only for testing purposes and is not currently planned for full support.

Bugzilla:1696451

Setting up a Samba server on an IdM domain member is provided 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

9.5. HIGH AVAILABILITY AND CLUSTERS

Pacemaker podman bundles available as a Technology Preview

Pacemaker container bundles now run on Podman, 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.

Bugzilla: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.

[Bugzilla:1784200](#)

New fence-agents-heuristics-ping fence agent

As a Technology Preview, Pacemaker now provides 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.

Bugzilla:1775847

9.6. 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 a Technology Preview.

Previously, the IdM API was enhanced to enable multiple versions of API commands. These 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 can 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 \(TECHNOLOGY PREVIEW\)](#).

Bugzilla:1664719

DNSSEC available as Technology Preview in IdM

Identity Management (IdM) servers with integrated DNS now implement 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](#)
- [Secure Domain Name System \(DNS\) Deployment Guide](#)
- [DNSSEC Key Rollover Timing Considerations](#)


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.

Bugzilla:1664718

ACME available as a Technology Preview

The Automated Certificate Management Environment (ACME) service is now available in Identity Management (IdM) as a Technology Preview. ACME is a protocol for automated identifier validation and certificate issuance. Its goal is to improve security by reducing certificate lifetimes and avoiding manual processes from certificate lifecycle management.

In RHEL, the ACME service uses the Red Hat Certificate System (RHCS) PKI ACME responder. The RHCS ACME subsystem is automatically deployed on every certificate authority (CA) server in the IdM deployment, but it does not service requests until the administrator enables it. RHCS uses the **acmeIPAServerCert** profile when issuing ACME certificates. The validity period of issued certificates is 90 days. Enabling or disabling the ACME service affects the entire IdM deployment.



IMPORTANT

It is recommended to enable ACME only in an IdM deployment where all servers are running RHEL 8.4 or later. Earlier RHEL versions do not include the ACME service, which can cause problems in mixed-version deployments. For example, a CA server without ACME can cause client connections to fail, because it uses a different DNS Subject Alternative Name (SAN).



WARNING

Currently, RHCS does not remove expired certificates. Because ACME certificates expire after 90 days, the expired certificates can accumulate and this can affect performance.

- To enable ACME across the whole IdM deployment, use the **ipa-acme-manage enable** command:

```
# ipa-acme-manage enable
The ipa-acme-manage command was successful
```

- To disable ACME across the whole IdM deployment, use the **ipa-acme-manage disable** command:

```
# ipa-acme-manage disable
The ipa-acme-manage command was successful
```

- To check whether the ACME service is installed and if it is enabled or disabled, use the **ipa-acme-manage status** command:

```
# ipa-acme-manage status
ACME is enabled
The ipa-acme-manage command was successful
```

Bugzilla:1628987

sssd-idp sub-package available as a Technology Preview

The **sssd-idp** sub-package for SSSD contains the **oidc_child** and **krb5 idp** plugins, which are client-side components that perform OAuth2 authentication against Identity Management (IdM) servers. This feature is available only with IdM servers on RHEL 8.7 and later.

[Bugzilla:2065692](#)

SSSD internal krb5 idp plugin available as a Technology Preview

The SSSD **krb5 idp** plugin allows you to authenticate against an external identity provider (IdP) using the OAuth2 protocol. This feature is available only with IdM servers on RHEL 8.7 and later.

[Bugzilla:2056483](#)

RHEL IdM allows delegating user authentication to external identity providers as a Technology Preview

As a Technology Preview in RHEL IdM, you can now associate users with external identity providers (IdP) that support the OAuth 2 device authorization flow. When these users authenticate with the SSSD version available in RHEL 8.7 or later, they receive RHEL IdM single sign-on capabilities with Kerberos tickets after performing authentication and authorization at the external IdP.

Notable features include:

- Adding, modifying, and deleting references to external IdPs with **ipa idp-*** commands
- Enabling IdP authentication for users with the **ipa user-mod --user-auth-type=idp** command

For additional information, see [Using external identity providers to authenticate to IdM](#) .

[Bugzilla:2101770](#)

9.7. DESKTOP

GNOME for the 64-bit ARM architecture available as a Technology Preview

The GNOME desktop environment is available for the 64-bit ARM architecture as a Technology Preview.

You can now connect to the desktop session on a 64-bit ARM server using VNC. As a result, you can manage the server using graphical applications.

A limited set of graphical applications is available on 64-bit ARM. For example:

- The Firefox web browser
- Red Hat Subscription Manager (**subscription-manager-cockpit**)
- Firewall Configuration (**firewall-config**)
- Disk Usage Analyzer (**baobab**)

Using Firefox, you can connect to the Cockpit service on the server.

Certain applications, such as LibreOffice, only provide a command-line interface, and their graphical interface is disabled.

Jira:RHELPLAN-27394, [Bugzilla:1667225](#), [Bugzilla:1724302](#), [Bugzilla:1667516](#)

GNOME for the IBM Z architecture available as a Technology Preview

The GNOME desktop environment is available for the IBM Z architecture as a Technology Preview.

You can now connect to the desktop session on an IBM Z server using VNC. As a result, you can manage the server using graphical applications.

A limited set of graphical applications is available on IBM Z. For example:

- The Firefox web browser
- Red Hat Subscription Manager (**subscription-manager-cockpit**)
- Firewall Configuration (**firewall-config**)
- Disk Usage Analyzer (**baobab**)

Using Firefox, you can connect to the Cockpit service on the server.

Certain applications, such as LibreOffice, only provide a command-line interface, and their graphical interface is disabled.

Jira:RHELPLAN-27737

9.8. 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.

Bugzilla:1698565

Intel Arc A-Series graphics available as a Technology Preview

Intel Arc A-Series graphics, also known as Alchemist or DG2, are now available as a Technology Preview.

To enable hardware acceleration with Intel Arc A-Series graphics, add the following option on the kernel command line:

```
i915.force_probe=pci-id
```

In this option, replace ***pci-id*** with either of the following:

- The PCI ID of your Intel GPU.
- The * character to enable the i915 driver with all alpha-quality hardware.

Bugzilla:2041686

9.9. 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 and AMD 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>

Bugzilla:1519039

AMD SEV and SEV-ES for KVM virtual machines

As a Technology Preview, RHEL 8 provides 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 the VM's memory to protect the VM from access by the host. This increases the security of the VM.

In addition, the enhanced Encrypted State version of SEV (SEV-ES) is also provided as Technology Preview. SEV-ES encrypts all CPU register contents when a VM stops running. This prevents the host from modifying the VM's CPU registers or reading any information from them.

Note that SEV and SEV-ES work only on the 2nd generation of AMD EPYC CPUs (codenamed Rome) or later. Also note that RHEL 8 includes SEV and SEV-ES encryption, but not the SEV and SEV-ES security attestation.

Bugzilla:1501618, Bugzilla: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, it is possible to enable a VNC console operated by Intel vGPU. By enabling it, users can connect to a VNC console of the VM and see the VM's desktop hosted by Intel vGPU. However, this currently only works for RHEL guest operating systems.

Bugzilla:1528684

Creating nested virtual machines

Nested KVM virtualization is provided as a Technology Preview for KVM virtual machines (VMs) running on Intel, AMD64, IBM POWER, 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.

Jira:RHELPLAN-14047, Jira:RHELPLAN-24437

Technology Preview: Select Intel network adapters now provide 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 that are supported by the **ixgbevf** and **iaavf** 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 provided with Microsoft Windows Server 2016 and later.

Bugzilla:1348508

Intel TDX in RHEL guests

As a Technology Preview, the Intel Trust Domain Extension (TDX) feature can now be used in RHEL 8.8 guest operating systems. If the host system supports TDX, you can deploy hardware-isolated RHEL 9 virtual machines (VMs), called trust domains (TDs). Note, however, that TDX currently does not work with **kdump**, and enabling TDX will cause **kdump** to fail on the VM.

Bugzilla:1836977

Sharing files between hosts and VMs using virtiofs

As a Technology Preview, RHEL 8 now provides the virtio file system (**virtiofs**). Using **virtiofs**, you can efficiently share files between your host system and its virtual machines (VM).

Bugzilla:1741615

9.10. RHEL IN CLOUD ENVIRONMENTS

RHEL confidential VMs are now available on Azure as a Technology Preview

With the updated RHEL kernel, you can now create and run confidential virtual machines (VMs) on Microsoft Azure as a Technology Preview. However, it is not yet possible to encrypt RHEL confidential VM images during boot on Azure.

Jira:RHELPLAN-122316

9.11. CONTAINERS

Clients for sigstore signatures with Fulcio and Rekor are now available as a Technology Preview

With Fulcio and Rekor servers, you can now create signatures by using short-term certificates based on an OpenID Connect (OIDC) server authentication, instead of manually managing a private key. Clients for sigstore signatures with Fulcio and Rekor are now available as a Technology Preview. This added functionality is the client side support only, and does not include either the Fulcio or Rekor servers.

Add the **fulcio** section in the **policy.json** file. To sign container images, use the **podman push --sign-by-sigstore=file.yml** or **skopeo copy --sign-by-sigstore=file.yml** commands, where **file.yml** is the sigstore signing parameter file.

To verify signatures, add the **fulcio** section and the **rekorPublicKeyPath** or **rekorPublicKeyData** fields in the **policy.json** file. For more information, see **containers-policy.json** man page.

Jira:RHELPLAN-136610

Quadlet in Podman is now available as a Technology Preview

Beginning with Podman v4.4, you can use Quadlet to automatically generate a **systemd** service file from the container description as a Technology Preview. The container description is in the **systemd** unit file format. The description focuses on the relevant container details and hides the technical complexity of running containers under **systemd**. The Quadlets are easier to write and maintain than the **systemd** unit files.

For more details, see the [upstream documentation](#) and [Make systemd better for Podman with Quadlet](#).

Jira:RHELPLAN-148394

The `podman-machine` command is unsupported

The **podman-machine** command for managing virtual machines, is available only as a Technology Preview. Instead, run Podman directly from the command line.

Jira:RHELDOCS-16861

CHAPTER 10. 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](#).

10.1. 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.

Bugzilla: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.

Bugzilla:1637872

The Kickstart `autostep` command has been deprecated

The `autostep` command has been deprecated. The related section about this command has been removed from the [RHEL 8 documentation](#).

Bugzilla:1904251

10.2. SUBSCRIPTION MANAGEMENT

The `--token` option of the `subscription-manager` command is deprecated

The `--token=<TOKEN>` option of the `subscription-manager register` command is an authentication method that helps register your system to Red Hat. This option depends on capabilities offered by the entitlement server. The default entitlement server, [subscription.rhsm.redhat.com](#), is planning to turn off this capability. As a consequence, attempting to use `subscription-manager register --token=<TOKEN>` might fail with the following error message:

```
Token authentication not supported by the entitlement server
```

You can continue registering your system using other authorization methods, such as including paired options `--username / --password` and `--org / --activationkey` of the `subscription-manager register` command.

Bugzilla:2170082

10.3. SOFTWARE MANAGEMENT

`rpmbuild --sign` is deprecated

The `rpmbuild --sign` command is deprecated since RHEL 8.1. 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.

Bugzilla:1688849

10.4. SHELLS AND COMMAND-LINE TOOLS

The `OpenEXR` component has been deprecated

The `OpenEXR` component has been deprecated. Hence, the support for the `EXR` image format has been dropped from the `imagecodecs` module.

[Bugzilla:1886310](#)

The **dump** utility from the **dump** package has been deprecated

The **dump** utility used for backup of file systems has been deprecated and will not be available in RHEL 9.

In RHEL 9, Red Hat recommends using the **tar**, **dd**, or **bacula**, backup utility, based on type of usage, which provides full and safe backups on ext2, ext3, and ext4 file systems.

Note that the **restore** utility from the **dump** package remains available and supported in RHEL 9 and is available as the **restore** package.

[Bugzilla:1997366](#)

The **hidepid=n** mount option is not supported in RHEL 8 **systemd**

The mount option **hidepid=n**, which controls who can access information in **/proc/[pid]** directories, is not compatible with **systemd** infrastructure provided in RHEL 8.

In addition, using this option might cause certain services started by **systemd** to produce SELinux AVC denial messages and prevent other operations from completing.

For more information, see the related Knowledgebase solution [Is mounting /proc with "hidepid=2" recommended with RHEL7 and RHEL8?](#)

[Bugzilla:2038929](#)

The **/usr/lib/udev/rename_device** utility has been deprecated

The **udev** helper utility **/usr/lib/udev/rename_device** for renaming network interfaces has been deprecated.

[Bugzilla:1875485](#)

The **ABRT** tool has been deprecated

The Automatic Bug Reporting Tool (ABRT) for detecting and reporting application crashes has been deprecated in RHEL 8. As a replacement, use the **systemd-coredump** tool to log and store core dumps, which are automatically generated files after a program crashes.

[Bugzilla:2055826](#)

The **ReaR** crontab has been deprecated

The **/etc/cron.d/rear** crontab from the **rear** package has been deprecated in RHEL 8 and will not be available in RHEL 9. The crontab checks every night whether the disk layout has changed, and runs **rear mkrescue** command if a change happened.

If you require this functionality, after an upgrade to RHEL 9, configure periodic runs of ReaR manually.

[Bugzilla:2083301](#)

The **SQLite** database backend in **Bacula** has been deprecated

The Bacula backup system supported multiple database backends: PostgreSQL, MySQL, and SQLite. The SQLite backend has been deprecated and will become unsupported in a later release of RHEL. As a replacement, migrate to one of the other backends (PostgreSQL or MySQL) and do not use the SQLite backend in new deployments.

[Bugzilla:2089399](#)

The **raw** command has been deprecated

The **raw** (`/usr/bin/raw`) command has been deprecated. Using this command in future releases of Red Hat Enterprise Linux can result in an error.

Jira:RHELPLAN-133171

10.5. 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.

[Bugzilla: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.

[Bugzilla: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.

[Bugzilla:1646541](#)

fapolicyd.rules is deprecated

The `/etc/fapolicyd/rules.d/` directory for files containing allow and deny execution rules replaces the `/etc/fapolicyd/fapolicyd.rules` file. The **fagenrules** script now merges all component rule files in this directory to the `/etc/fapolicyd/compiled.rules` file. Rules in `/etc/fapolicyd/fapolicyd.trust` are still processed by the **fapolicyd** framework but only for ensuring backward compatibility.

[Bugzilla:2054741](#)

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.

Bugzilla:1645153

Runtime disabling SELinux using `/etc/selinux/config` is now deprecated

Runtime disabling SELinux using the **SELINUX=disabled** option in the `/etc/selinux/config` file has been deprecated. In RHEL 9, when you disable SELinux only through `/etc/selinux/config`, the system starts with SELinux enabled but with no policy loaded.

If your scenario really requires to completely disable SELinux, Red Hat recommends disabling 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.

Bugzilla:1932222

The `ipa` SELinux module removed from `selinux-policy`

The `ipa` SELinux module has been removed from the `selinux-policy` package because it is no longer maintained. The functionality is now included in the `ipa-selinux` subpackage.

If your scenario requires the use of types or interfaces from the `ipa` module in a local SELinux policy, install the `ipa-selinux` package.

Bugzilla:1461914

TPM 1.2 is deprecated

The Trusted Platform Module (TPM) secure cryptoprocessor standard 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.

Bugzilla:1657927

crypto-policies derived properties are now deprecated

With the introduction of scopes for **crypto-policies** directives in custom policies, the following derived properties have been deprecated: **tls_cipher**, **ssh_cipher**, **ssh_group**, **ike_protocol**, and **sha1_in_dnssec**. Additionally, the use of the **protocol** property without specifying a scope is now deprecated as well. See the **crypto-policies(7)** man page for recommended replacements.

Bugzilla:2011208

10.6. 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.

Bugzilla:1647725

The dropwatch tool is deprecated

The **dropwatch** tool has been deprecated. The tool will not be supported in future releases, thus it is not recommended for new deployments. As a replacement of this package, Red Hat recommends to use the **perf** command line tool.

For more information on using the **perf** command line tool, see the [Getting started with Perf](#) section on the Red Hat customer portal or the **perf** man page.

Bugzilla:1929173

The xinetd service has been deprecated

The **xinetd** service has been deprecated and will be removed in RHEL 9. As a replacement, use **systemd**. For further details, see [How to convert xinetd service to systemd](#) .

Bugzilla:2009113

The cgdcbxd package is deprecated

Control group data center bridging exchange daemon (**cgdcbxd**) is a service to monitor data center bridging (DCB) netlink events and manage the **net_prio_control** group subsystem. Starting with RHEL 8.5, the **cgdcbxd** package is deprecated and will be removed in the next major RHEL release.

Bugzilla:2006665

The WEP Wi-Fi connection method is deprecated

The insecure wired equivalent privacy (WEP) Wi-Fi connection method is deprecated in RHEL 8 and will be removed in RHEL 9.0. For secure Wi-Fi connections, use the Wi-Fi Protected Access 3 (WPA3) or WPA2 connection methods.

Bugzilla:2029338

The unsupported xt_u32 module is now deprecated

Using the unsupported **xt_u32** module, users of **iptables** can match arbitrary 32 bits in the packet header or payload. Since RHEL 8.6, the **xt_u32** module is deprecated and will be removed in RHEL 9.

If you use **xt_u32**, migrate to the **nftables** packet filtering framework. For example, first change your firewall to use **iptables** with native matches to incrementally replace individual rules, and later use the **iptables-translate** and accompanying utilities to migrate to **nftables**. If no native match exists in **nftables**, use the raw payload matching feature of **nftables**. For details, see the **raw payload expression** section in the **nft(8)** man page.

Bugzilla:2061288

The term **slaves** is deprecated in the **nmstate** API

Red Hat is committed to using conscious language. See details about this initiative in [Making open source more inclusive](#). Therefore the **slaves** term is deprecated in the Nmstate API. Use the term **port** when you use **nmstatectl**.

(Jira:RHELDPCS-17641)

10.7. KERNEL

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.

Bugzilla:1878207

The Linux **firewire** sub-system and its associated user-space components are deprecated in RHEL 8

The **firewire** sub-system provides interfaces to use and maintain any resources on the IEEE 1394 bus. In RHEL 9, **firewire** will no longer be supported in the **kernel** package. Note that **firewire** contains several user-space components provided by the **libavc1394**, **libdc1394**, **libraw1394** packages. These packages are subject to the deprecation as well.

Bugzilla:1871863

Installing RHEL for Real Time 8 using diskless boot is now deprecated

Diskless booting allows multiple systems to share a root file system through the network. While convenient, diskless boot is prone to introducing network latency in real-time workloads. With a future minor update of RHEL for Real Time 8, the diskless booting feature will no longer be supported.

[Bugzilla:1748980](#)

Kernel live patching now covers all RHEL minor releases

Since RHEL 8.1, kernel live patches have been provided for selected minor release streams of RHEL covered under the Extended Update Support (EUS) policy to remediate Critical and Important Common Vulnerabilities and Exposures (CVEs). To accommodate the maximum number of concurrently covered kernels and use cases, the support window for each live patch has been decreased from 12 to 6 months for every minor, major, and zStream version of the kernel. It means that on the day a kernel live patch is released, it will cover every minor release and scheduled errata kernel delivered in the past 6 months.

For more information about this feature, see [Applying patches with kernel live patching](#).

For details about available kernel live patches, see [Kernel Live Patch life cycles](#).

[Bugzilla:1958250](#)

The **crash-ptdump-command** package is deprecated

The **crash-ptdump-command** package, which is a **ptdump** extension module for the crash utility, is deprecated and might not be available in future RHEL releases. The **ptdump** command fails to retrieve the log buffer when working in the Single Range Output mode and only works in the Table of Physical Addresses (ToPA) mode. **crash-ptdump-command** is currently not maintained upstream

Bugzilla:1838927

10.8. BOOT LOADER

The **kernelopts** environment variable has been deprecated

In RHEL 8, the kernel command-line parameters for systems using the GRUB bootloader were defined in the **kernelopts** environment variable. The variable was stored in the `/boot/grub2/grubenv` file for each kernel boot entry. However, storing the kernel command-line parameters using **kernelopts** was not robust. Therefore, with a future major update of RHEL, **kernelopts** will be removed and the kernel command-line parameters will be stored in the Boot Loader Specification (BLS) snippet instead.

[Bugzilla:2060759](#)

10.9. 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](#).

Bugzilla:1665295

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.

Bugzilla:1592011

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.

[Bugzilla:1871953](#)

VDO write modes other than **async** are deprecated

VDO supports several write modes in RHEL 8:

- **sync**

- **async**
- **async-unsafe**
- **auto**

Starting with RHEL 8.4, the following write modes are deprecated:

sync

Devices above the VDO layer cannot recognize if VDO is synchronous, and consequently, the devices cannot take advantage of the VDO **sync** mode.

async-unsafe

VDO added this write mode as a workaround for the reduced performance of **async** mode, which complies to Atomicity, Consistency, Isolation, and Durability (ACID). Red Hat does not recommend **async-unsafe** for most use cases and is not aware of any users who rely on it.

auto

This write mode only selects one of the other write modes. It is no longer necessary when VDO supports only a single write mode.

These write modes will be removed in a future major RHEL release.

The recommended VDO write mode is now **async**.

For more information on VDO write modes, see [Selecting a VDO write mode](#).

Jira:RHELPLAN-70700

VDO manager has been deprecated

The python-based VDO management software has been deprecated and will be removed from RHEL 9. In RHEL 9, it will be replaced by the LVM-VDO integration. Therefore, it is recommended to create VDO volumes using the **lvcreate** command.

The existing volumes created using the VDO management software can be converted using the **/usr/sbin/lvm_import_vdo** script, provided by the **lvm2** package. For more information on the LVM-VDO implementation, see [Deduplicating and compressing logical volumes on RHEL](#).

[Bugzilla:1949163](#)

cramfs has been deprecated

Due to lack of users, the **cramfs** kernel module is deprecated. **squashfs** is recommended as an alternative solution.

[Bugzilla:1794513](#)

10.10. HIGH AVAILABILITY AND CLUSTERS

pcs commands that support the clufter tool have been deprecated

The **pcs** commands that support the **clufter** tool for analyzing cluster configuration formats have been deprecated. These commands now print a warning that the command has been deprecated and sections related to these commands have been removed from the **pcs** help display and the **pcs(8)** man page.

The following commands have been deprecated:

- **pcs config import-cman** for importing CMAN / RHEL6 HA cluster configuration
- **pcs config export** for exporting cluster configuration to a list of **pcs** commands which recreate the same cluster

Bugzilla:1851335

10.11. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

The **mod_php** module provided with PHP for use with the Apache HTTP Server has been deprecated

The **mod_php** module provided with PHP for use with the Apache HTTP Server in RHEL 8 is available but not enabled in the default configuration. The module is no longer available in RHEL 9.

Since RHEL 8, PHP scripts are run using the FastCGI Process Manager (**php-fpm**) by default. For more information, see [Using PHP with the Apache HTTP Server](#).

Bugzilla:2225332

10.12. COMPILERS AND DEVELOPMENT TOOLS

The **gdb.i686** packages are deprecated

In RHEL 8.1, the 32-bit versions of the GNU Debugger (GDB), **gdb.i686**, were shipped due to a dependency problem in another package. Because RHEL 8 does not support 32-bit hardware, the **gdb.i686** packages are deprecated since RHEL 8.4. The 64-bit versions of GDB, **gdb.x86_64**, are fully capable of debugging 32-bit applications.

If you use **gdb.i686**, note the following important issues:

- The **gdb.i686** packages will no longer be updated. Users must install **gdb.x86_64** instead.
- If you have **gdb.i686** installed, installing **gdb.x86_64** will cause **yum** to report **package gdb-8.2-14.el8.x86_64 obsoletes gdb < 8.2-14.el8 provided by gdb-8.2-12.el8.i686**. This is expected. Either uninstall **gdb.i686** or pass **dnf** the **--allowerase** option to remove **gdb.i686** and install **gdb.x86_64**.
- Users will no longer be able to install the **gdb.i686** packages on 64-bit systems, that is, those with the **libc.so.6()(64-bit)** packages.

Bugzilla:1853140

libdwarf has been deprecated

The **libdwarf** library has been deprecated in RHEL 8. The library will likely not be supported in future major releases. Instead, use the **elfutils** and **libdw** libraries for applications that wish to process ELF/DWARF files.

Alternatives for the **libdwarf-tools dwarfdump** program are the **binutils readelf** program or the **elfutils eu-readelf** program, both used by passing the **--debug-dump** flag.

Bugzilla:1920624

10.13. 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.

[Bugzilla: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.

[Bugzilla:1877991](#)

The SSSD version of `libwbclient` has been removed

The SSSD implementation of the `libwbclient` package was deprecated in RHEL 8.4. As it cannot be used with recent versions of Samba, the SSSD implementation of `libwbclient` has now been removed.

[Bugzilla:1947671](#)

Standalone use of the `ctdb` service has been deprecated

Since RHEL 8.4, customers are advised to use the `ctdb` clustered Samba service only when both of the following conditions apply:

- The `ctdb` service is managed as a **pacemaker** resource with the resource-agent `ctdb`.
- The `ctdb` service uses storage volumes that contain either a GlusterFS file system provided by the Red Hat Gluster Storage product or a GFS2 file system.

The stand-alone use case of the `ctdb` service has been deprecated and will not be included in a next major release of Red Hat Enterprise Linux. For further information on support policies for Samba, see the Knowledgebase article [Support Policies for RHEL Resilient Storage - ctdb General Policies](#) .

[Bugzilla:1916296](#)

Indirect AD integration with IdM via WinSync has been deprecated

WinSync is no longer actively developed in RHEL 8 due to several functional limitations:

- WinSync supports only one Active Directory (AD) domain.
- Password synchronization requires installing additional software on AD Domain Controllers.

For a more robust solution with better resource and security separation, Red Hat recommends using a **cross-forest trust** for indirect integration with Active Directory. See the [Indirect integration](#) documentation.

[Jira:RHELPLAN-100400](#)

The SSSD implicit files provider domain is disabled by default

The default value of the **enable_files_domain** setting in the `/etc/sss/sss.conf` configuration file has been changed from **true** to **false**. This means that the SSSD implicit **files** provider domain, which retrieves user and group information from local files `/etc/passwd` and `/etc/group`, is now disabled by default.

The default **glibc files** module, instead of SSSD, serves local users. SSSD does not start automatically, unless you have defined a domain in the `sss.conf` file.

The implementation of the SSSD **files** provider is still available for explicit configuration for specific use cases, such as smart card authentication of local users.

Jira:RHELPLAN-139456

Running Samba as a PDC or BDC is deprecated

The classic domain controller mode that enabled administrators to run Samba as an NT4-like primary domain controller (PDC) and backup domain controller (BDC) is deprecated. The code and settings to configure these modes will be removed in a future Samba release.

As long as the Samba version in RHEL 8 provides the PDC and BDC modes, Red Hat supports these modes only in existing installations with Windows versions which support NT4 domains. Red Hat recommends not setting up a new Samba NT4 domain, because Microsoft operating systems later than Windows 7 and Windows Server 2008 R2 do not support NT4 domains.

If you use the PDC to authenticate only Linux users, Red Hat suggests migrating to [Red Hat Identity Management \(IdM\)](#) that is included in RHEL subscriptions. However, you cannot join Windows systems to an IdM domain. Note that Red Hat continues supporting the PDC functionality IdM uses in the background.

Red Hat does not support running Samba as an AD domain controller (DC).

[Bugzilla:1926114](#)

The SMB1 protocol is deprecated in Samba

Starting with Samba 4.11, the insecure Server Message Block version 1 (SMB1) protocol is deprecated and will be removed in a future release.


To improve the security, by default, SMB1 is disabled in the Samba server and client utilities.

Jira:RHELDOCS-16612

Limited support for FreeRADIUS

In RHEL 8, the following external authentication modules are deprecated as part of the FreeRADIUS offering:

- The MySQL, PostgreSQL, SQLite, and unixODBC database connectors
- The **Perl** language module
- The REST API module



NOTE

The PAM authentication module and other authentication modules that are provided as part of the base package are not affected.

You can find replacements for the deprecated modules in community-supported packages, for example in the Fedora project.

In addition, the scope of support for the **freeradius** package will be limited to the following use cases in future RHEL releases:

- Using FreeRADIUS as a wireless-authentication provider with Identity Management (IdM) as the backend source of authentication. The authentication occurs through the **krb5** and LDAP authentication packages or as PAM authentication in the main FreeRADIUS package.

- Using FreeRADIUS to provide a source-of-truth for authentication in IdM, through the Python 3 authentication package.

In contrast to these deprecations, Red Hat will strengthen the support of the following external authentication modules with FreeRADIUS:

- Authentication based on **krb5** and LDAP
- **Python 3** authentication

The focus on these integration options is in close alignment with the strategic direction of Red Hat IdM.

Jira:RHELDOCS-17573

10.14. 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.

Bugzilla:1607766

10.15. 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.

Bugzilla:1569610

Motif has been deprecated

The Motif widget toolkit has been deprecated in RHEL, because development in the upstream Motif community is inactive.

The following Motif packages have been deprecated, including their development and debugging variants:

- **motif**
- **openmotif**
- **openmotif21**
- **openmotif22**

Additionally, the **motif-static** package has been removed.

Red Hat recommends using the GTK toolkit as a replacement. GTK is more maintainable and provides new features compared to Motif.

Jira:RHELPLAN-98983

10.16. 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.

[Bugzilla:1666722](#)

10.17. 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.

[Bugzilla:1874892](#)

The `network` System Role displays a deprecation warning when configuring teams on RHEL 9 nodes

The network teaming capabilities have been deprecated in RHEL 9. As a result, using the `network` RHEL System Role on an RHEL 8 control node to configure a network team on RHEL 9 nodes, shows a warning about the deprecation.

[Bugzilla:2021685](#)

Ansible Engine has been deprecated

Previous versions of RHEL 8 provided access to an Ansible Engine repository, with a limited scope of support, to enable supported RHEL Automation use cases, such as RHEL System Roles and Insights remediations. Ansible Engine has been deprecated, and Ansible Engine 2.9 will have no support after September 29, 2023. For more details on the supported use cases, see [Scope of support for the Ansible Core package included in the RHEL 9 AppStream](#).

Users must manually migrate their systems from Ansible Engine to Ansible Core. For that, follow the steps:

Procedure

1. Check if the system is running RHEL 8.7 or a later release:

```
# cat /etc/redhat-release
```

2. Uninstall Ansible Engine 2.9:

```
# yum remove ansible
```

3. Disable the `ansible-2-for-rhel-8-x86_64-rpms` repository:

```
# subscription-manager repos --disable  
ansible-2-for-rhel-8-x86_64-rpms
```

4. Install the Ansible Core package from the RHEL 8 AppStream repository:

```
# yum install ansible-core
```

For more details, see: [Using Ansible in RHEL 8.6 and later](#) .

[Bugzilla:2006081](#)

10.18. VIRTUALIZATION

virsh iface-* commands have become deprecated

The **virsh iface-*** commands, such as **virsh iface-start** and **virsh iface-destroy**, are now deprecated, and will be removed in a future major version of RHEL. In addition, these commands 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, such as **nmcli**.

[Bugzilla:1664592](#)

virt-manager has been deprecated

The Virtual Machine Manager application, also known as **virt-manager**, has been deprecated. The RHEL 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 in the RHEL web console.

[Jira:RHELPLAN-10304](#)

Limited support for virtual machine snapshots

Creating snapshots of virtual machines (VMs) is currently only supported for VMs not using the UEFI firmware. In addition, during the snapshot operation, the QEMU monitor may become blocked, which negatively impacts the hypervisor performance for certain workloads.

Also note that the current mechanism of creating VM snapshots has been deprecated, and Red Hat does not recommend using VM snapshots in a production environment.

[Bugzilla: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**.

[Bugzilla:1651994](#)

SPICE has been deprecated

The SPICE remote display protocol has become 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.

Bugzilla:1849563

KVM on IBM POWER has been deprecated

Using KVM virtualization on IBM POWER hardware has become deprecated. As a result, KVM on IBM POWER is still supported in RHEL 8, but will become unsupported in a future major release of RHEL.

Jira:RHELPLAN-71200

SecureBoot image verification using SHA1-based signatures is deprecated

Performing SecureBoot image verification using SHA1-based signatures on UEFI (PE/COFF) executables has become deprecated. Instead, Red Hat recommends using signatures based on the SHA2 algorithm, or later.

Bugzilla:1935497

Using SPICE to attach smart card readers to virtual machines has been deprecated

The SPICE remote display protocol has been deprecated in RHEL 8. Since the only recommended way to attach smart card readers to virtual machines (VMs) depends on the SPICE protocol, the usage of smart cards in VMs has also become deprecated in RHEL 8.

In a future major version of RHEL, the functionality of attaching smart card readers to VMs will only be supported by third party remote visualization solutions.

[Bugzilla:2059626](#)

RDMA-based live migration is deprecated

With this update, migrating running virtual machines using Remote Direct Memory Access (RDMA) has become deprecated. As a result, it is still possible to use the **rdma://** migration URI to request migration over RDMA, but this feature will become unsupported in a future major release of RHEL.

Jira:RHELPLAN-153267

10.19. CONTAINERS

The Podman varlink-based API v1.0 has been removed

The Podman varlink-based API v1.0 was deprecated in a previous release of RHEL 8. Podman v2.0 introduced a new Podman v2.0 RESTful API. With the release of Podman v3.0, the varlink-based API v1.0 has been completely removed.

Jira:RHELPLAN-45858

container-tools:1.0 has been deprecated

The **container-tools:1.0** module has been deprecated and will no longer receive security updates. It is recommended to use a newer supported stable module stream, such as **container-tools:2.0** or **container-tools:3.0**.

Jira:RHELPLAN-59825

The container-tools:2.0 module has been deprecated

The `container-tools:2.0` module has been deprecated and will no longer receive security updates. It is recommended to use a newer supported stable module stream, such as **container-tools:3.0**.

Jira:RHELPLAN-85066

Flatpak images except GIMP has been deprecated

The **rhel8/firefox-flatpak**, **rhel8/thunderbird-flatpak**, **rhel8/inkscape-flatpak**, and **rhel8/libreoffice-flatpak** RHEL 8 Flatpak Applications have been deprecated and replaced by the RHEL 9 versions. The **rhel8/gimp-flatpak** Flatpak Application is not deprecated because there is no replacement yet in RHEL 9.

[Bugzilla:2142499](#)

The CNI network stack has been deprecated

The Container Network Interface (CNI) network stack will be deprecated in a future minor version. Previously, containers connected to the single Container Network Interface (CNI) plugin only via DNS. Podman v.4.0 introduced a new Netavark network stack. You can use the Netavark network stack with Podman and other Open Container Initiative (OCI) container management applications. The Netavark network stack for Podman is also compatible with advanced Docker functionalities. Containers in multiple networks can access containers on any of those networks.

For more information, see [Switching the network stack from CNI to Netavark](#) .

Jira:RHELPLAN-145958

container-tools:3.0 has been deprecated

The **container-tools:3.0** module has been deprecated and will no longer receive security updates. To continue to build and run Linux Containers on RHEL, use a newer, stable, and supported module stream, such as **container-tools:4.0**.


For instructions on switching to a later stream, see [Switching to a later stream](#) .

Jira:RHELPLAN-146398

10.20. DEPRECATED PACKAGES

This section lists packages that have been deprecated and will probably not be included in a future major release of Red Hat Enterprise Linux.

For changes to packages between RHEL 7 and RHEL 8, see [Changes to packages](#) in the *Considerations in adopting RHEL 8* document.



IMPORTANT

The support status of deprecated packages remains unchanged within RHEL 8. For more information about the length of support, see [Red Hat Enterprise Linux Life Cycle](#) and [Red Hat Enterprise Linux Application Streams Life Cycle](#) .

The following packages have been deprecated in RHEL 8:

- `389-ds-base-legacy-tools`
- `abrt`

- `abrt-addon-ccpp`
- `abrt-addon-kerneloops`
- `abrt-addon-pstoreoops`
- `abrt-addon-vmcore`
- `abrt-addon-xorg`
- `abrt-cli`
- `abrt-console-notification`
- `abrt-dbus`
- `abrt-desktop`
- `abrt-gui`
- `abrt-gui-libs`
- `abrt-libs`
- `abrt-tui`
- `adobe-source-sans-pro-fonts`
- `adwaita-qt`
- `alsa-plugins-pulseaudio`
- `amanda`
- `amanda-client`
- `amanda-libs`
- `amanda-server`
- `ant-contrib`
- `antlr3`
- `antlr32`
- `aopalliance`
- `apache-commons-collections`
- `apache-commons-compress`
- `apache-commons-exec`
- `apache-commons-jxpath`
- `apache-commons-parent`

- apache-ivy
- apache-parent
- apache-resource-bundles
- apache-sshd
- apiguardian
- aspnetcore-runtime-3.0
- aspnetcore-runtime-3.1
- aspnetcore-runtime-5.0
- aspnetcore-targeting-pack-3.0
- aspnetcore-targeting-pack-3.1
- aspnetcore-targeting-pack-5.0
- assertj-core
- authd
- auto
- autoconf213
- autogen
- autogen-libopts
- awscli
- base64coder
- batik
- batik-css
- batik-util
- bea-stax
- bea-stax-api
- bind-export-devel
- bind-export-libs
- bind-libs-lite
- bind-pkcs11
- bind-pkcs11-devel

- bind-pkcs11-libs
- bind-pkcs11-utils
- bind-sdb
- bind-sdb
- bind-sdb-chroot
- bluez-hid2hci
- boost-jam
- boost-signals
- bouncycastle
- bpg-algeti-fonts
- bpg-chveulebrivi-fonts
- bpg-classic-fonts
- bpg-courier-fonts
- bpg-courier-s-fonts
- bpg-dedaena-block-fonts
- bpg-dejavu-sans-fonts
- bpg-elite-fonts
- bpg-excelsior-caps-fonts
- bpg-excelsior-condenced-fonts
- bpg-excelsior-fonts
- bpg-fonts-common
- bpg-glaho-fonts
- bpg-gorda-fonts
- bpg-ingiri-fonts
- bpg-irubaqidze-fonts
- bpg-mikhail-stephan-fonts
- bpg-mrgvlovani-caps-fonts
- bpg-mrgvlovani-fonts
- bpg-nateli-caps-fonts

- `bpg-nateli-condenced-fonts`
- `bpg-nateli-fonts`
- `bpg-nino-medium-cond-fonts`
- `bpg-nino-medium-fonts`
- `bpg-sans-fonts`
- `bpg-sans-medium-fonts`
- `bpg-sans-modern-fonts`
- `bpg-sans-regular-fonts`
- `bpg-serif-fonts`
- `bpg-serif-modern-fonts`
- `bpg-ucnobi-fonts`
- `brlapi-java`
- `bsh`
- `buildnumber-maven-plugin`
- `byaccj`
- `call0n`
- `cbi-plugins`
- `cdparanoia`
- `cdparanoia-devel`
- `cdparanoia-libs`
- `cdrdao`
- `cmirror`
- `codehaus-parent`
- `codemodel`
- `compat-exiv2-026`
- `compat-guile18`
- `compat-hwloc1`
- `compat-libpthread-nonshared`
- `compat-libtiff3`

- compat-openssl10
- compat-sap-c++-11
- compat-sap-c++-10
- compat-sap-c++-9
- createrepo_c-devel
- ctags
- ctags-etags
- custodia
- cyrus-imapd-vzic
- dbus-c++
- dbus-c++-devel
- dbus-c++-glib
- dbxtool
- dhcp-libs
- directory-maven-plugin
- directory-maven-plugin-javadoc
- dirsplit
- dleyna-connector-dbus
- dleyna-core
- dleyna-renderer
- dleyna-server
- dnssec-trigger
- dnssec-trigger-panel
- dotnet-apphost-pack-3.0
- dotnet-apphost-pack-3.1
- dotnet-apphost-pack-5.0
- dotnet-host-fxr-2.1
- dotnet-host-fxr-2.1
- dotnet-hostfxr-3.0

- dotnet-hostfxr-3.1
- dotnet-hostfxr-5.0
- dotnet-runtime-2.1
- dotnet-runtime-3.0
- dotnet-runtime-3.1
- dotnet-runtime-5.0
- dotnet-sdk-2.1
- dotnet-sdk-2.1.5xx
- dotnet-sdk-3.0
- dotnet-sdk-3.1
- dotnet-sdk-5.0
- dotnet-targeting-pack-3.0
- dotnet-targeting-pack-3.1
- dotnet-targeting-pack-5.0
- dotnet-templates-3.0
- dotnet-templates-3.1
- dotnet-templates-5.0
- dotnet5.0-build-reference-packages
- dptfextract
- drpm
- drpm-devel
- dump
- dvd+rw-tools
- dyninst-static
- eclipse-ecf
- eclipse-ecf-core
- eclipse-ecf-runtime
- eclipse-emf
- eclipse-emf-core

- eclipse-emf-runtime
- eclipse-emf-xsd
- eclipse-equinox-osgi
- eclipse-jdt
- eclipse-license
- eclipse-p2-discovery
- eclipse-pde
- eclipse-platform
- eclipse-swt
- ed25519-java
- ee4j-parent
- elfutils-devel-static
- elfutils-libelf-devel-static
- enca
- enca-devel
- environment-modules-compat
- evince-browser-plugin
- exec-maven-plugin
- farstream02
- felix-gogo-command
- felix-gogo-runtime
- felix-gogo-shell
- felix-scr
- felix-osgi-compendium
- felix-osgi-core
- felix-osgi-foundation
- felix-parent
- file-roller
- fipscheck

- fipscheck-devel
- fipscheck-lib
- firewire
- fonts-tweak-tool
- forge-parent
- freeradius-mysql
- freeradius-perl
- freeradius-postgresql
- freeradius-rest
- freeradius-sqlite
- freeradius-unixODBC
- fuse-sshfs
- fusesource-pom
- future
- gamin
- gamin-devel
- gavl
- gcc-toolset-10
- gcc-toolset-10-annobin
- gcc-toolset-10-binutils
- gcc-toolset-10-binutils-devel
- gcc-toolset-10-build
- gcc-toolset-10-dwz
- gcc-toolset-10-dyninst
- gcc-toolset-10-dyninst-devel
- gcc-toolset-10-elfutils
- gcc-toolset-10-elfutils-debuginfod-client
- gcc-toolset-10-elfutils-debuginfod-client-devel
- gcc-toolset-10-elfutils-devel

- gcc-toolset-10-elfutils-libelf
- gcc-toolset-10-elfutils-libelf-devel
- gcc-toolset-10-elfutils-libs
- gcc-toolset-10-gcc
- gcc-toolset-10-gcc-c++
- gcc-toolset-10-gcc-gdb-plugin
- gcc-toolset-10-gcc-gfortran
- gcc-toolset-10-gdb
- gcc-toolset-10-gdb-doc
- gcc-toolset-10-gdb-gdbserver
- gcc-toolset-10-libasan-devel
- gcc-toolset-10-libatomic-devel
- gcc-toolset-10-libitm-devel
- gcc-toolset-10-libsan-devel
- gcc-toolset-10-libquadmath-devel
- gcc-toolset-10-libstdc++-devel
- gcc-toolset-10-libstdc++-docs
- gcc-toolset-10-libtsan-devel
- gcc-toolset-10-libubsan-devel
- gcc-toolset-10-ltrace
- gcc-toolset-10-make
- gcc-toolset-10-make-devel
- gcc-toolset-10-perftools
- gcc-toolset-10-runtime
- gcc-toolset-10-strace
- gcc-toolset-10-systemtap
- gcc-toolset-10-systemtap-client
- gcc-toolset-10-systemtap-devel
- gcc-toolset-10-systemtap-initscript

- gcc-toolset-10-systemtap-runtime
- gcc-toolset-10-systemtap-sdt-devel
- gcc-toolset-10-systemtap-server
- gcc-toolset-10-toolchain
- gcc-toolset-10-valgrind
- gcc-toolset-10-valgrind-devel
- gcc-toolset-9
- gcc-toolset-9-annobin
- gcc-toolset-9-build
- gcc-toolset-9-perftools
- gcc-toolset-9-runtime
- gcc-toolset-9-toolchain
- gcc-toolset-11-make-devel
- GConf2
- GConf2-devel
- gegl
- genisoimage
- genwqe-tools
- genwqe-vpd
- genwqe-zlib
- genwqe-zlib-devel
- geoipupdate
- geronimo-annotation
- geronimo-jms
- geronimo-jpa
- geronimo-parent-poms
- gfbgraph
- gflags
- gflags-devel

- glassfish-annotation-api
- glassfish-el
- glassfish-fastinfoset
- glassfish-jaxb-core
- glassfish-jaxb-txw2
- glassfish-jsp
- glassfish-jsp-api
- glassfish-legal
- glassfish-master-pom
- glassfish-servlet-api
- glew-devel
- glib2-fam
- glog
- glog-devel
- gmock
- gmock-devel
- gnome-abrt
- gnome-boxes
- gnome-menus-devel
- gnome-online-miners
- gnome-shell-extension-disable-screenshield
- gnome-shell-extension-horizontal-workspaces
- gnome-shell-extension-no-hot-corner
- gnome-shell-extension-window-grouper
- gnome-themes-standard
- gnu-free-fonts-common
- gnu-free-mono-fonts
- gnu-free-sans-fonts
- gnu-free-serif-fonts

- gnupg2-smime
- gnuplot
- gnuplot-common
- gobject-introspection-devel
- google-gson
- google-noto-sans-syriac-eastern-fonts
- google-noto-sans-syriac-estrangela-fonts
- google-noto-sans-syriac-western-fonts
- google-noto-sans-tibetan-fonts
- google-noto-sans-ui-fonts
- gphoto2
- gsl-devel
- gssntlmssp
- gtest
- gtest-devel
- gtkmm24
- gtkmm24-devel
- gtkmm24-docs
- gtksourceview3
- gtksourceview3-devel
- gtkspell
- gtkspell-devel
- gtkspell3
- guile
- gutenprint-gimp
- gutenprint-libs-ui
- gvfs-afc
- gvfs-afp
- gvfs-archive

- hamcrest-core
- hawtjni
- hawtjni
- hawtjni-runtime
- HdrHistogram
- HdrHistogram-javadoc
- highlight-gui
- hivex-devel
- hostname
- hplip-gui
- httpcomponents-project
- hwloc-plugins
- hyphen-fo
- hyphen-grc
- hyphen-hsb
- hyphen-ia
- hyphen-is
- hyphen-ku
- hyphen-mi
- hyphen-mn
- hyphen-sa
- hyphen-tk
- ibus-sayura
- icedax
- icu4j
- idm-console-framework
- inkscape
- inkscape-docs
- inkscape-view

- iptables
- ipython
- isl
- isl-devel
- isorelax
- istack-commons-runtime
- istack-commons-tools
- iwl3945-firmware
- iwl4965-firmware
- iwl6000-firmware
- jacoco
- jaf
- jaf-javadoc
- jakarta-oro
- janino
- jansi-native
- jarjar
- java-1.8.0-ibm
- java-1.8.0-ibm-demo
- java-1.8.0-ibm-devel
- java-1.8.0-ibm-headless
- java-1.8.0-ibm-jdbc
- java-1.8.0-ibm-plugin
- java-1.8.0-ibm-src
- java-1.8.0-ibm-webstart
- java-1.8.0-openjdk-accessibility
- java-1.8.0-openjdk-accessibility-slowdebug
- java_cup
- java-atk-wrapper

- javacc
- javacc-maven-plugin
- javaewah
- javaparser
- javapoet
- javassist
- javassist-javadoc
- jaxen
- jboss-annotations-1.2-api
- jboss-interceptors-1.2-api
- jboss-logmanager
- jboss-parent
- jctools
- jdepend
- jdependency
- jdom
- jdom2
- jetty
- jetty-continuation
- jetty-http
- jetty-io
- jetty-security
- jetty-server
- jetty-servlet
- jetty-util
- jffi
- jflex
- jgit
- jline

- jmc
- jnr-netdb
- jolokia-jvm-agent
- js-uglify
- jsch
- json_simple
- jss-javadoc
- jtidy
- junit5
- jvnet-parent
- jzlib
- kernel-cross-headers
- ksc
- kurdit-unikurd-web-fonts
- kyotocabinet-libs
- ldapjdk-javadoc
- lensfun
- lensfun-devel
- lftp-scripts
- libaec
- libaec-devel
- libappindicator-gtk3
- libappindicator-gtk3-devel
- libatomic-static
- libavc1394
- libblocksruntime
- libcacard
- libcacard-devel
- libcgroup

- libcgroup-tools
- libchamplain
- libchamplain-devel
- libchamplain-gtk
- libcroco
- libcroco-devel
- libcxl
- libcxl-devel
- libdap
- libdap-devel
- libdazzle-devel
- libdbusmenu
- libdbusmenu-devel
- libdbusmenu-doc
- libdbusmenu-gtk3
- libdbusmenu-gtk3-devel
- libdc1394
- libdnet
- libdnet-devel
- libdv
- libdwarf
- libdwarf-devel
- libdwarf-static
- libdwarf-tools
- libeasyfc
- libeasyfc-gobject
- libepubgen-devel
- libertas-sd8686-firmware
- libertas-usb8388-firmware

- `libertas-usb8388-olpc-firmware`
- `libgdither`
- `libGLEW`
- `libgovirt`
- `libguestfs-benchmarking`
- `libguestfs-devel`
- `libguestfs-gfs2`
- `libguestfs-gobject`
- `libguestfs-gobject-devel`
- `libguestfs-java`
- `libguestfs-java-devel`
- `libguestfs-javadoc`
- `libguestfs-man-pages-ja`
- `libguestfs-man-pages-uk`
- `libguestfs-tools`
- `libguestfs-tools-c`
- `libhugetlbfs`
- `libhugetlbfs-devel`
- `libhugetlbfs-utils`
- `libIDL`
- `libIDL-devel`
- `libidn`
- `libiec61883`
- `libindicator-gtk3`
- `libindicator-gtk3-devel`
- `libiscsi-devel`
- `libjose-devel`
- `libkkc`
- `libkkc-common`

- libkkc-data
- libldb-devel
- liblogging
- libluksmeta-devel
- libmalaga
- libmcpp
- libmemcached
- libmemcached-libs
- libmetalink
- libmodulemd1
- libmongocrypt
- libmtp-devel
- libmusicbrainz5
- libmusicbrainz5-devel
- libnbd-devel
- liboauth
- liboauth-devel
- libpfm-static
- libpng12
- libpurple
- libpurple-devel
- libraw1394
- libreport-plugin-mailx
- libreport-plugin-rhtsupport
- libreport-plugin-ureport
- libreport-rhel
- libreport-rhel-bugzilla
- librpmem
- librpmem-debug

- librpmem-devel
- libsass
- libsass-devel
- libselinux-python
- libsqlite3x
- libtalloc-devel
- libtar
- libtdb-devel
- libtevent-devel
- libtpms-devel
- libunwind
- libusal
- libvarlink
- libverto-libevent
- libvirt-admin
- libvirt-bash-completion
- libvirt-daemon-driver-storage-gluster
- libvirt-daemon-driver-storage-iscsi-direct
- libvirt-devel
- libvirt-docs
- libvirt-gconfig
- libvirt-gobject
- libvirt-lock-sanlock
- libvirt-wireshark
- libvmem
- libvmem-debug
- libvmem-devel
- libvmmalloc
- libvmmalloc-debug

- libvmmalloc-devel
- libvncserver
- libwinpr-devel
- libwmf
- libwmf-devel
- libwmf-lite
- libXNVCtrl
- libyami
- log4j12
- log4j12-javadoc
- lohit-malayalam-fonts
- lohit-nepali-fonts
- lorax-composer
- lua-guestfs
- lucene
- lucene-analysis
- lucene-analyzers-smartcn
- lucene-queries
- lucene-queryparser
- lucene-sandbox
- lz4-java
- lz4-java-javadoc
- mailman
- mailx
- make-devel
- malaga
- malaga-suomi-voikko
- marisa
- maven-antrun-plugin

- maven-assembly-plugin
- maven-clean-plugin
- maven-dependency-analyzer
- maven-dependency-plugin
- maven-doxia
- maven-doxia-sitetools
- maven-install-plugin
- maven-invoker
- maven-invoker-plugin
- maven-parent
- maven-plugins-pom
- maven-reporting-api
- maven-reporting-impl
- maven-resolver-api
- maven-resolver-connector-basic
- maven-resolver-impl
- maven-resolver-spi
- maven-resolver-transport-wagon
- maven-resolver-util
- maven-scm
- maven-script-interpreter
- maven-shade-plugin
- maven-shared
- maven-verifier
- maven-wagon-file
- maven-wagon-http
- maven-wagon-http-shared
- maven-wagon-provider-api
- maven2

- meanwhile
- mercurial
- mercurial-hgk
- metis
- metis-devel
- mingw32-bzip2
- mingw32-bzip2-static
- mingw32-cairo
- mingw32-expat
- mingw32-fontconfig
- mingw32-freetype
- mingw32-freetype-static
- mingw32-gstreamer1
- mingw32-harfbuzz
- mingw32-harfbuzz-static
- mingw32-icu
- mingw32-libjpeg-turbo
- mingw32-libjpeg-turbo-static
- mingw32-libpng
- mingw32-libpng-static
- mingw32-libtiff
- mingw32-libtiff-static
- mingw32-openssl
- mingw32-readline
- mingw32-sqlite
- mingw32-sqlite-static
- mingw64-adwaita-icon-theme
- mingw64-bzip2
- mingw64-bzip2-static

- mingw64-cairo
- mingw64-expat
- mingw64-fontconfig
- mingw64-freetype
- mingw64-freetype-static
- mingw64-gstreamer1
- mingw64-harfbuzz
- mingw64-harfbuzz-static
- mingw64-icu
- mingw64-libjpeg-turbo
- mingw64-libjpeg-turbo-static
- mingw64-libpng
- mingw64-libpng-static
- mingw64-libtiff
- mingw64-libtiff-static
- mingw64-nettle
- mingw64-openssl
- mingw64-readline
- mingw64-sqlite
- mingw64-sqlite-static
- modello
- mojo-parent
- mongo-c-driver
- mousetweaks
- mozjs52
- mozjs52-devel
- mozjs60
- mozjs60-devel
- mozvoikko

- msv-javadoc
- msv-manual
- munge-maven-plugin
- mythes-mi
- mythes-ne
- nafees-web-naskh-fonts
- nbd
- nbdkit-devel
- nbdkit-example-plugins
- nbdkit-gzip-plugin
- nbdkit-plugin-python-common
- nbdkit-plugin-vddk
- ncompress
- ncurses-compat-libs
- net-tools
- netcf
- netcf-devel
- netcf-libs
- network-scripts
- network-scripts-ppp
- nkf
- nodejs-devel
- nodejs-packaging
- nss_nis
- nss-pam-ldapd
- objectweb-asm
- objectweb-asm-javadoc
- objectweb-pom
- ocaml-bisect-ppx

- ocaml-camlp4
- ocaml-camlp4-devel
- ocaml-lwt
- ocaml-mmap
- ocaml-ocplib-endian
- ocaml-ounit
- ocaml-result
- ocaml-seq
- opencryptoki-tpmtok
- opencv-contrib
- opencv-core
- opencv-devel
- openhpi
- openhpi-libs
- OpenIPMI-perl
- openssh-cavs
- openssh-ldap
- openssl-ibmpkcs11
- opentest4j
- os-maven-plugin
- pakchois
- pandoc
- paps-libs
- paranamer
- parfait
- parfait-examples
- parfait-javadoc
- pcp-parfait-agent
- pcp-pmda-rpm

- pcp-pmda-vmware
- pcsc-lite-doc
- peripety
- perl-B-Debug
- perl-B-Lint
- perl-Class-Factory-Util
- perl-Class-ISA
- perl-DateTime-Format-HTTP
- perl-DateTime-Format-Mail
- perl-File-CheckTree
- perl-homedir
- perl-libxml-perl
- perl-Locale-Codes
- perl-Mozilla-LDAP
- perl-NKF
- perl-Object-HashBase-tools
- perl-Package-DeprecationManager
- perl-Pod-LaTeX
- perl-Pod-Plainer
- perl-prefork
- perl-String-CRC32
- perl-SUPER
- perl-Sys-Virt
- perl-tests
- perl-YAML-Syck
- phodav
- php-recode
- php-xmlrpc
- pidgin

- pidgin-devel
- pidgin-sipe
- pinentry-emacs
- pinentry-gtk
- pipewire0.2-devel
- pipewire0.2-libs
- platform-python-coverage
- plexus-ant-factory
- plexus-bsh-factory
- plexus-cli
- plexus-component-api
- plexus-component-factories-pom
- plexus-components-pom
- plexus-i18n
- plexus-interactivity
- plexus-pom
- plexus-velocity
- plymouth-plugin-throbgress
- pmreorder
- postgresql-test-rpm-macros
- powermock
- prometheus-jmx-exporter
- prometheus-jmx-exporter-openjdk11
- ptscotch-mpich
- ptscotch-mpich-devel
- ptscotch-mpich-devel-parmetis
- ptscotch-openmpi
- ptscotch-openmpi-devel
- purple-sipe

- pygobject2-doc
- pygtk2
- pygtk2-codegen
- pygtk2-devel
- pygtk2-doc
- python-nose-docs
- python-nss-doc
- python-podman-api
- python-psycopg2-doc
- python-pymongo-doc
- python-redis
- python-schedutils
- python-slip
- python-sqlalchemy-doc
- python-varlink
- python-virtualenv-doc
- python2-backports
- python2-backports-ssl_match_hostname
- python2-bson
- python2-coverage
- python2-docs
- python2-docs-info
- python2-funcsigs
- python2-ipaddress
- python2-mock
- python2-nose
- python2-numpy-doc
- python2-psycopg2-debug
- python2-psycopg2-tests

- python2-pymongo
- python2-pymongo-gridfs
- python2-pytest-mock
- python2-sqlalchemy
- python2-tools
- python2-virtualenv
- python3-bson
- python3-click
- python3-coverage
- python3-cpio
- python3-custodia
- python3-docs
- python3-flask
- python3-gevent
- python3-gobject-base
- python3-hivex
- python3-html5lib
- python3-hypothesis
- python3-ipatests
- python3-itsdangerous
- python3-jwt
- python3-libguestfs
- python3-mock
- python3-networkx-core
- python3-nose
- python3-nss
- python3-openipmi
- python3-pillow
- python3-ptyprocess

- python3-pydbus
- python3-pymongo
- python3-pymongo-gridfs
- python3-pyOpenSSL
- python3-pytoml
- python3-reportlab
- python3-schedutils
- python3-scons
- python3-semantic_version
- python3-slip
- python3-slip-dbus
- python3-sqlalchemy
- python3-syspurpose
- python3-virtualenv
- python3-webencodings
- python3-werkzeug
- python38-asn1crypto
- python38-numpy-doc
- python38-psycopg2-doc
- python38-psycopg2-tests
- python39-numpy-doc
- python39-psycopg2-doc
- python39-psycopg2-tests
- qemu-kvm-block-gluster
- qemu-kvm-block-iscsi
- qemu-kvm-block-ssh
- qemu-kvm-hw-usbredir
- qemu-kvm-device-display-virtio-gpu-gl
- qemu-kvm-device-display-virtio-gpu-pci-gl

- `qemu-kvm-device-display-virtio-vga-gl`
- `qemu-kvm-tests`
- `qpdf`
- `qpdf-doc`
- `qpidd-proton`
- `qrencode`
- `qrencode-devel`
- `qrencode-libs`
- `qt5-qtcanvas3d`
- `qt5-qtcanvas3d-examples`
- `rarian`
- `rarian-compat`
- `re2c`
- `recode`
- `redhat-lsb`
- `redhat-lsb-core`
- `redhat-lsb-cxx`
- `redhat-lsb-desktop`
- `redhat-lsb-languages`
- `redhat-lsb-printing`
- `redhat-lsb-submod-multimedia`
- `redhat-lsb-submod-security`
- `redhat-lsb-supplemental`
- `redhat-lsb-trialuse`
- `redhat-menus`
- `redhat-support-lib-python`
- `redhat-support-tool`
- `reflections`
- `regexp`

- relaxngDatatype
- rasm-gtk
- rpm-plugin-priorset
- rpemd
- rsyslog-udpspoof
- ruby-hivex
- ruby-libguestfs
- rubygem-abrt
- rubygem-abrt-doc
- rubygem-bson
- rubygem-bson-doc
- rubygem-bundler-doc
- rubygem-mongo
- rubygem-mongo-doc
- rubygem-net-telnet
- rubygem-xmlrpc
- s390utils-cmsfs
- samba-pidl
- samba-test
- samba-test-libs
- samyak-devanagari-fonts
- samyak-fonts-common
- samyak-gujarati-fonts
- samyak-malayalam-fonts
- samyak-odia-fonts
- samyak-tamil-fonts
- sane-frontends
- sanlk-reset
- sat4j

- scala
- scotch
- scotch-devel
- SDL_sound
- selinux-policy-minimum
- sendmail
- sgabios
- sgabios-bin
- shrinkwrap
- sisu-inject
- sisu-mojos
- sisu-plexus
- skkdic
- SLOF
- smc-anjalioldlipi-fonts
- smc-dyuthi-fonts
- smc-fonts-common
- smc-kalyani-fonts
- smc-raghmalayalam-fonts
- smc-suruma-fonts
- softhsm-devel
- sonatype-oss-parent
- sonatype-plugins-parent
- sos-collector
- sparsehash-devel
- spax
- spec-version-maven-plugin
- spice
- spice-client-win-x64

- spice-client-win-x86
- spice-glib
- spice-glib-devel
- spice-gtk
- spice-gtk-tools
- spice-gtk3
- spice-gtk3-devel
- spice-gtk3-vala
- spice-parent
- spice-protocol
- spice-qxl-wddm-dod
- spice-server
- spice-server-devel
- spice-qxl-xddm
- spice-server
- spice-streaming-agent
- spice-vdagent-win-x64
- spice-vdagent-win-x86
- sssd-libwbclient
- star
- stax-ex
- stax2-api
- stringtemplate
- stringtemplate4
- subscription-manager-initial-setup-addon
- subscription-manager-migration
- subscription-manager-migration-data
- subversion-javahl
- SuperLU

- SuperLU-devel
- supermin-devel
- swig
- swig-doc
- swig-gdb
- swtpm-devel
- swtpm-tools-pkcs11
- system-storage-manager
- tcl-brlapi
- testng
- tibetan-machine-uni-fonts
- timedatex
- tpm-quote-tools
- tpm-tools
- tpm-tools-pkcs11
- treelayout
- trousers
- trousers-lib
- tuned-profiles-compatible
- tuned-profiles-nfv-host-bin
- tuned-utils-systemtap
- tycho
- uglify-js
- unbound-devel
- univocity-output-tester
- univocity-parsers
- usbguard-notifier
- usbredir-devel
- utf8cpp

- uthash
- velocity
- vinagre
- vino
- virt-dib
- virt-p2v-maker
- vm-dump-metrics-devel
- weld-parent
- wodim
- woodstox-core
- wqy-microhei-fonts
- wqy-unibit-fonts
- xdelta
- xmlgraphics-commons
- xmlstreambuffer
- xinetd
- xorg-x11-apps
- xorg-x11-drv-qxl
- xorg-x11-server-Xspice
- xpp3
- xsane-gimp
- xsom
- xz-java
- xz-java-javadoc
- yajl-devel
- yp-tools
- ypbind
- ypserv

10.21. DEPRECATED AND UNMAINTAINED DEVICES

This section lists devices (drivers, adapters) that

- continue to be supported until the end of life of RHEL 8 but will likely not be supported in future major releases of this product and are not recommended for new deployments. Support for devices other than those listed remains unchanged. These are **deprecated** devices.
- are available but are no longer being tested or updated on a routine basis in RHEL 8. Red Hat may fix serious bugs, including security bugs, at its discretion. These devices should no longer be used in production, and it is likely they will be disabled in the next major release. These are **unmaintained** devices.

PCI device IDs are in the format of *vendor:device:subvendor:subdevice*. If no device ID is listed, all devices associated with the corresponding driver have been deprecated. To check the PCI IDs of the hardware on your system, run the **lspci -nn** command.

Table 10.1. Deprecated devices

Device ID	Driver	Device name
	bnx2	QLogic BCM5706/5708/5709/5716 Driver
	hpsa	Hewlett-Packard Company: Smart Array Controllers
0x10df:0x0724	lpfc	Emulex Corporation: OneConnect FCoE Initiator (Skyhawk)
0x10df:0xe200	lpfc	Emulex Corporation: LPe15000/LPe16000 Series 8Gb/16Gb Fibre Channel Adapter
0x10df:0xf011	lpfc	Emulex Corporation: Saturn: LightPulse Fibre Channel Host Adapter
0x10df:0xf015	lpfc	Emulex Corporation: Saturn: LightPulse Fibre Channel Host Adapter
0x10df:0xf100	lpfc	Emulex Corporation: LPe12000 Series 8Gb Fibre Channel Adapter
0x10df:0xfc40	lpfc	Emulex Corporation: Saturn-X: LightPulse Fibre Channel Host Adapter
0x10df:0xe220	be2net	Emulex Corporation: OneConnect NIC (Lancer)
0x1000:0x005b	megaraid_sas	Broadcom / LSI: MegaRAID SAS 2208 [Thunderbolt]
0x1000:0x006E	mpt3sas	Broadcom / LSI: SAS2308 PCI-Express Fusion-MPT SAS-2
0x1000:0x0080	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0081	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2

Device ID	Driver	Device name
0x1000:0x0082	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0083	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0084	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0085	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0086	mpt3sas	Broadcom / LSI: SAS2308 PCI-Express Fusion-MPT SAS-2
0x1000:0x0087	mpt3sas	Broadcom / LSI: SAS2308 PCI-Express Fusion-MPT SAS-2
	myri10g e	Myricom 10G driver (10GbE)
	netxen_ nic	QLogic/NetXen (1/10) GbE Intelligent Ethernet Driver
0x1077:0x2031	qla2xxx	QLogic Corp.: ISP8324-based 16Gb Fibre Channel to PCI Express Adapter
0x1077:0x2532	qla2xxx	QLogic Corp.: ISP2532-based 8Gb Fibre Channel to PCI Express HBA
0x1077:0x8031	qla2xxx	QLogic Corp.: 8300 Series 10GbE Converged Network Adapter (FCoE)
	qla3xxx	QLogic ISP3XXX Network Driver v2.03.00-k5
0x1924:0x0803	sfc	Solarflare Communications: SFC9020 10G Ethernet Controller
0x1924:0x0813	sfc	Solarflare Communications: SFL9021 10GBASE-T Ethernet Controller
	Soft- RoCE (rdma_ r xe)	
	HNS- RoCE	HNS GE/10GE/25GE/50GE/100GE RDMA Network Controller
	liquidio	Cavium LiquidIO Intelligent Server Adapter Driver

Device ID	Driver	Device name
	liquidio_vf	Cavium LiquidIO Intelligent Server Adapter Virtual Function Driver

Table 10.2. Unmaintained devices

Device ID	Driver	Device name
	e1000	Intel® PRO/1000 Network Driver
	mptbase	Fusion MPT SAS Host driver
	mptsas	Fusion MPT SAS Host driver
	mptscsih	Fusion MPT SCSI Host driver
	mptspi	Fusion MPT SAS Host driver
0x1000:0x0071 ^[a]	megaraid_sas	Broadcom / LSI: MR SAS HBA 2004
0x1000:0x0073 ^[a]	megaraid_sas	Broadcom / LSI: MegaRAID SAS 2008 [Falcon]
0x1000:0x0079 ^[a]	megaraid_sas	Broadcom / LSI: MegaRAID SAS 2108 [Liberator]
	nvmet_tcp	NVMe/TCP target driver
^[a] Disabled in RHEL 8.0, re-enabled in RHEL 8.4 due to customer requests.		

CHAPTER 11. KNOWN ISSUES

This part describes known issues in Red Hat Enterprise Linux 8.8.

11.1. INSTALLER AND IMAGE CREATION

Installation fails on IBM Power 10 systems with LPAR and secure boot enabled

RHEL installer is not integrated with static key secure boot on IBM Power 10 systems. Consequently, when logical partition (LPAR) is enabled with the secure boot option, the installation fails with the error, **Unable to proceed with RHEL-x.x Installation**.

To work around this problem, install RHEL without enabling secure boot. After booting the system:

1. Copy the signed Kernel into the PReP partition using the **dd** command.
2. Restart the system and enable secure boot.

Once the firmware verifies the bootloader and the kernel, the system boots up successfully.

For more information, see <https://www.ibm.com/support/pages/node/6528884>

Bugzilla:2025814

Unexpected SELinux policies on systems where Anaconda is running as an application

When Anaconda is running as an application on an already installed system (for example to perform another installation to an image file using the **-image** anaconda option), the system is not prohibited to modify the SELinux types and attributes during installation. As a consequence, certain elements of SELinux policy might change on the system where Anaconda is running. To work around this problem, do not run Anaconda on the production system and execute it in a temporary virtual machine. So that the SELinux policy on a production system is not modified. Running anaconda as part of the system installation process such as installing from **boot.iso** or **dvd.iso** is not affected by this issue.

Bugzilla:2050140

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.

Bugzilla: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.

Bugzilla:1697896

The USB CD-ROM drive is not available as an installation source in Anaconda

Installation fails when the USB CD-ROM drive is the source for it and the Kickstart **ignoredisk --only-use=** command is specified. In this case, Anaconda cannot find and use this source disk.

To work around this problem, use the **harddrive --partition=sdX --dir=/** command to install from USB CD-ROM drive. As a result, the installation does not fail.

[Bugzilla:1914955](#)

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.

[Bugzilla:1757877](#)

Hard drive partitioned installations with iso9660 filesystem fails

You cannot install RHEL on systems where the hard drive is partitioned with the **iso9660** filesystem. This is due to the updated installation code that is set to ignore any hard disk containing a **iso9660** file system partition. This happens even when RHEL is installed without using a DVD.

To work around this problem, add the following script in the kickstart file to format the disc before the installation starts.

Note: Before performing the workaround, backup the data available on the disk. The **wipefs** command formats all the existing data from the disk.

```
%pre
wipefs -a /dev/sda
%end
```

As a result, installations work as expected without any errors.

[Bugzilla:1929105](#)

IBM Power systems with HASH MMU mode fail to boot with memory allocation failures

IBM Power Systems with **HASH memory allocation unit (MMU)** mode support **kdump** up to a maximum of 192 cores. Consequently, the system fails to boot with memory allocation failures if **kdump** is enabled on more than 192 cores. This limitation is due to RMA memory allocations during early boot in **HASH MMU** mode. To work around this problem, use the **Radix MMU** mode with **fadump** enabled instead of using **kdump**.

[Bugzilla:2028361](#)

RHEL for Edge installer image fails to create mount points when installing an rpm-ostree payload

When deploying **rpm-ostree** payloads, used for example in a RHEL for Edge installer image, the installer does not properly create some mount points for custom partitions. As a consequence, the installation is aborted with the following error:

The command 'mount --bind /mnt/sysimage/data /mnt/sysroot/data' exited with the code 32.

To work around this issue:

- Use an automatic partitioning scheme and do not add any mount points manually.
- Manually assign mount points only inside **/var** directory. For example, **/var/my-mount-point**, and the following standard directories: **/**, **/boot**, **/var**.

As a result, the installation process finishes successfully.

[Bugzilla:2126506](#)

11.2. SUBSCRIPTION MANAGEMENT

syspurpose addons 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 **addons**. 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 **addons** argument will not observe any effect on the subscriptions that are auto-attached.

[Bugzilla:1687900](#)

11.3. SOFTWARE MANAGEMENT

cr_compress_file_with_stat() can cause a memory leak

The **createrepo_c** C library has the API **cr_compress_file_with_stat()** function. This function is declared with **char **dst** as a second parameter. Depending on its other parameters, **cr_compress_file_with_stat()** either uses **dst** as an input parameter, or uses it to return an allocated string. This unpredictable behavior can cause a memory leak, because it does not inform the user when to free **dst** contents.

To work around this problem, a new API **cr_compress_file_with_stat_v2** function has been added, which uses the **dst** parameter only as an input. It is declared as **char *dst**. This prevents memory leak.

Note that the **cr_compress_file_with_stat_v2** function is temporary and will be present only in RHEL 8. Later, **cr_compress_file_with_stat()** will be fixed instead.

[Bugzilla:1973588](#)

YUM transactions reported as successful when a scriptlet fails

Since RPM version 4.6, post-install scriptlets are allowed to fail without being fatal to the transaction. This behavior propagates up to YUM as well. This results in scriptlets which might occasionally fail while the overall package transaction reports as successful.

There is no workaround available at the moment.

Note that this is expected behavior that remains consistent between RPM and YUM. Any issues in scriptlets should be addressed at the package level.

[Bugzilla:1986657](#)

11.4. SHELLS AND COMMAND-LINE TOOLS

ipmitool is incompatible with certain server platforms

The **ipmitool** utility serves for monitoring, configuring, and managing devices that support the Intelligent Platform Management Interface (IPMI). The current version of **ipmitool** uses Cipher Suite 17 by default instead of the previous Cipher Suite 3. Consequently, **ipmitool** fails to communicate with certain bare metal nodes that announced support for Cipher Suite 17 during negotiation, but do not actually support this cipher suite. As a result, **ipmitool** aborts with the **no matching cipher suite** error message.

For more details, see the related [Knowledgebase article](#).

To solve this problem, update your baseboard management controller (BMC) firmware to use the Cipher Suite 17.

Optionally, if the BMC firmware update is not available, you can work around this problem by forcing **ipmitool** to use a certain cipher suite. When invoking a managing task with **ipmitool**, add the **-C** option to the **ipmitool** command together with the *number* of the cipher suite you want to use. See the following example:

```
# ipmitool -I lanplus -H myserver.example.com -P mypass -C 3 chassis power status
```

[Bugzilla:1873614](#)

ReaR fails to recreate a volume group when you do not use clean disks for restoring

ReaR fails to perform recovery when you want to restore to disks that contain existing data.

To work around this problem, wipe the disks manually before restoring to them if they have been previously used. To wipe the disks in the rescue environment, use one of the following commands before running the **rear recover** command:

- The **dd** command to overwrite the disks.
- The **wipefs** command with the **-a** flag to erase all available metadata.

See the following example of wiping metadata from the **/dev/sda** disk:

```
# wipefs -a /dev/sda[1-9] /dev/sda
```

This command wipes the metadata from the partitions on **/dev/sda** first, and then the partition table itself.

[Bugzilla:1925531](#)

coreutils might report misleading EPERM error codes

GNU Core Utilities (**coreutils**) started using the **statx()** system call. If a **seccomp** filter returns an EPERM error code for unknown system calls, **coreutils** might consequently report misleading EPERM error codes because EPERM can not be distinguished from the actual *Operation not permitted* error

returned by a working **statx()** syscall.

To work around this problem, update the **seccomp** filter to either permit the **statx()** syscall, or to return an ENOSYS error code for syscalls it does not know.

[Bugzilla:2030661](#)

11.5. INFRASTRUCTURE SERVICES

Postfix TLS fingerprint algorithm in the FIPS mode needs to be changed to SHA-256

By default in RHEL 8, **postfix** uses MD5 fingerprints with the TLS for backward compatibility. But in the FIPS mode, the MD5 hashing function is not available, which may cause TLS to incorrectly function in the default postfix configuration. To workaround this problem, the hashing function needs to be changed to SHA-256 in the postfix configuration file.

For more details, see the related Knowledgebase article [Fix postfix TLS in the FIPS mode by switching to SHA-256 instead of MD5](#).

[Bugzilla:1711885](#)

The **brlitty** package is not multilib compatible

It is not possible to have both 32-bit and 64-bit versions of the **brlitty** package installed. You can either install the 32-bit (**brlitty.i686**) or the 64-bit (**brlitty.x86_64**) version of the package. The 64-bit version is recommended.

[Bugzilla:2008197](#)

11.6. SECURITY

tangd-keygen does not handle non-default **umask** correctly

The **tangd-keygen** script does not change file permissions for generated key files. Consequently, on systems with a default user file-creation mode mask (**umask**) that prevents reading keys to other users, the **tang-show-keys** command returns the error message **Internal Error 500** instead of displaying the keys.

To work around the problem, use the **chmod o+r *.jwk** command to change permissions on the files in the **/var/db/tang** directory.

[Bugzilla:2188743](#)

sshd -T provides inaccurate information about Ciphers, MACs and KeX algorithms

The output of the **sshd -T** command does not contain the system-wide crypto policy configuration or other options that could come from an environment file in **/etc/sysconfig/ssh** and that are applied as arguments on the **sshd** command. This occurs because the upstream OpenSSH project did not support the Include directive to support Red-Hat-provided cryptographic defaults in RHEL 8. Crypto policies are applied as command-line arguments to the **sshd** executable in the **sshd.service** unit during the service's start by using an **EnvironmentFile**. To work around the problem, use the **source** command with the environment file and pass the crypto policy as an argument to the **sshd** command, as in **sshd -T \$CRYPTO_POLICY**. For additional information, see [Ciphers, MACs or KeX algorithms differ from sshd -T to what is provided by current crypto policy level](#). As a result, the output from **sshd -T** matches the currently configured crypto policy.

[Bugzilla:2044354](#)

RHV hypervisor may not work correctly when hardening the system during installation

When installing Red Hat Virtualization Hypervisor (RHV-H) and applying the Red Hat Enterprise Linux 8 STIG profile, OSCP Anaconda Add-on may harden the system as RHEL instead of RVH-H and remove essential packages for RHV-H. Consequently, the RHV hypervisor may not work. To work around the problem, install the RHV-H system without applying any profile hardening, and after the installation is complete, apply the profile by using OpenSCAP. As a result, the RHV hypervisor works correctly.

[Bugzilla:2075508](#)

CVE OVAL feeds are now only in the compressed format, and data streams are not in the SCAP 1.3 standard

Red Hat provides CVE OVAL feeds in the bzip2-compressed format and are no longer available in the XML file format. Because referencing compressed content is not standardized in the Security Content Automation Protocol (SCAP) 1.3 specification, third-party SCAP scanners can have problems scanning rules that use the feed.

[Bugzilla:2028428](#)

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 the Rsyslog remote logging application:

```
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.

[Bugzilla:1679512](#)

Server with GUI and Workstation installations are not possible with CIS Server profiles

The CIS Server Level 1 and Level 2 security profiles are 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 Server profiles is not possible. An attempted installation using the CIS Server Level 1 or Level 2 profiles 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.
```

If you need to align systems with the **Server with GUI** or **Workstation** software selections according to CIS benchmarks, use the CIS Workstation Level 1 or Level 2 profiles instead.

[Bugzilla:1843932](#)

Kickstart uses `org_fedora_oscaps` instead of `com_redhat_oscaps` in RHEL 8

The Kickstart references the Open Security Content Automation Protocol (OSCAP) Anaconda add-on as **org_fedora_oscaped** instead of **com_redhat_oscaped**, which might cause confusion. This is necessary to keep compatibility with Red Hat Enterprise Linux 7.

Bugzilla:1665082

libvirt overrides xccdf_org.ssgproject.content_rule_sysctl_net_ipv4_conf_all_forwarding

The **libvirt** virtualization framework enables IPv4 forwarding whenever a virtual network with a forward mode of **route** or **nat** is started. This overrides the configuration by the **xccdf_org.ssgproject.content_rule_sysctl_net_ipv4_conf_all_forwarding** rule, and subsequent compliance scans report the **fail** result when assessing this rule.

Apply one of these scenarios to work around the problem:

- Uninstall the **libvirt** packages if your scenario does not require them.
- Change the forwarding mode of virtual networks created by **libvirt**.
- Remove the **xccdf_org.ssgproject.content_rule_sysctl_net_ipv4_conf_all_forwarding** rule by tailoring your profile.

Bugzilla:2118758

OpenSSL in FIPS mode accepts only specific D-H parameters

In FIPS mode, 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 Diffie-Hellman 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.

Bugzilla:1810911

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.

Bugzilla:1919155

OpenSC might not detect CardOS V5.3 card objects correctly

The OpenSC toolkit does not correctly detect serial numbers of smart cards using the CardOS V5.3 system. Consequently, the **pkcs11-tool** utility might not list card objects.

To work around the problem, turn off file caching by setting the `use_file_caching = false` option in the `/etc/opensc.conf` file.

[Bugzilla:2176973](#)

Smart-card provisioning process through OpenSC `pkcs15-init` does not work properly

The `file_caching` option is enabled in the default OpenSC configuration, and the file caching functionality does not handle some commands from the `pkcs15-init` tool properly. Consequently, the smart-card provisioning process through OpenSC fails.

To work around the problem, add the following snippet to the `/etc/opensc.conf` file:

```
app pkcs15-init {
    framework pkcs15 {
        use_file_caching = false;
    }
}
```

The smart-card provisioning through `pkcs15-init` only works if you apply the previously described workaround.

[Bugzilla:1947025](#)

Connections to servers with SHA-1 signatures do not work with GnuTLS

SHA-1 signatures in certificates are rejected by the GnuTLS secure communications library as insecure. Consequently, applications that use GnuTLS as a TLS backend cannot establish a TLS connection to peers that offer such certificates. This behavior is inconsistent with other system cryptographic libraries.

To work around this problem, upgrade the server to use certificates signed with SHA-256 or stronger hash, or switch to the LEGACY policy.

[Bugzilla:1628553](#)

`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 `yum 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:

```
# yum module enable libselinux-python
# yum install libselinux-python
```

Alternatively, install `libselinux-python` using its install profile with a single command:

```
# yum module install libselinux-python:2.8/common
```

As a result, you can install `libselinux-python` using the respective module.

[Bugzilla: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.

[Bugzilla: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

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

IKE over TCP connections do not work on custom TCP ports

The **tcp-remoteport** Libreswan configuration option does not work properly. Consequently, an IKE over TCP connection cannot be established when a scenario requires specifying a non-default TCP port.

[Bugzilla:1989050](#)

scap-security-guide cannot configure termination of idle sessions

Even though the **sshd_set_idle_timeout** rule still exists in the data stream, the former method for idle session timeout of configuring **sshd** is no longer available. Therefore, the rule is marked as **not applicable** and cannot harden anything. Other methods for configuring idle session termination, such as **systemd** (Logind), are also not available. As a consequence, **scap-security-guide** cannot configure the system to reliably disconnect idle sessions after a certain amount of time.

You can work around this problem in one of the following ways, which might fulfill the security requirement:

- Configuring the **accounts_tmout** rule. However, this variable could be overridden by using the **exec** command.
- Configuring the **configure_tmux_lock_after_time** and **configure_bashrc_exec_tmux** rules. This requires installing the **tmux** package.
- Upgrading to RHEL 8.7 or later where the **systemd** feature is already implemented together with the proper SCAP rule.

[Bugzilla:2167373](#)

The OSCAP Anaconda add-on does not fetch tailored profiles in the graphical installation

The OSCAP Anaconda add-on does not provide an option to select or deselect tailoring of security profiles in the RHEL graphical installation. Starting from RHEL 8.8, the add-on does not take tailoring into account by default when installing from archives or RPM packages. Consequently, the installation displays the following error message instead of fetching an OSCAP tailored profile:

```
There was an unexpected problem with the supplied content.
```

To work around this problem, you must specify paths in the `%addon org_fedora_oscap` section of your Kickstart file, for example:

```
xccdf-path = /usr/share/xml/scap/sc_tailoring/ds-combined.xml
tailoring-path = /usr/share/xml/scap/sc_tailoring/tailoring-xccdf.xml
```

As a result, you can use the graphical installation for OSCAP tailored profiles only with the corresponding Kickstart specifications.

[Bugzilla:2165948](#)

The automatic screen lock does not work when a smart-card reader is removed

The `opensc` packages incorrectly handle removing USB smart-card readers. Consequently, the system remains unlocked even when the GNOME Display Manager (GDM) is configured to lock the screen when a smart card is removed. Furthermore, after you reconnect the USB reader, the screen also does not lock after removing the smart card.

To work around this problem, perform one of the following actions:

- Always remove only a smart card, not a smart-card reader.
- When using hardware tokens that integrate a reader and a card in one package, upgrade to RHEL 9.

[Bugzilla:2097048](#)

OpenSCAP memory-consumption problems

On systems with limited memory, the OpenSCAP scanner might terminate prematurely or it might not generate the results files. To work around this problem, you can customize the scanning profile to deselect rules that involve recursion over the entire / file system:

- `rpm_verify_hashes`
- `rpm_verify_permissions`
- `rpm_verify_ownership`
- `file_permissions_unauthorized_world_writable`
- `no_files_unowned_by_user`
- `dir_perms_world_writable_system_owned`
- `file_permissions_unauthorized_suid`

- `file_permissions_unauthorized_sgid`
- `file_permissions_ungroupowned`
- `dir_perms_world_writable_sticky_bits`

For more details and more workarounds, see the related [Knowledgebase article](#).

[Bugzilla:2161499](#)

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.

[Bugzilla:1834716](#)

11.7. NETWORKING

Systems with the `IPv6_rpfilter` option enabled experience low network throughput

Systems with the `IPv6_rpfilter` option enabled in the `firewalld.conf` file currently experience suboptimal performance and low network throughput in high traffic scenarios, such as 100 Gbps links. To work around the problem, disable the `IPv6_rpfilter` option. To do so, add the following line in the `/etc/firewalld/firewalld.conf` file.

```
IPv6_rpfilter=no
```

As a result, the system performs better, but also has reduced security.

[Bugzilla:1871860](#)

11.8. KERNEL

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 0x30000000-0x31fffff not reserved in ACPI namespace" appears during system boot](#) solution.

Bugzilla:1868526

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.

Bugzilla:1609288

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.

Bugzilla:1602962

Reloading an identical crash extension may cause segmentation faults

When you load a copy of an already loaded crash extension file, it might trigger a segmentation fault. Currently, the crash utility detects if an original file has been loaded. Consequently, due to two identical files co-existing in the crash utility, a namespace collision occurs, which triggers the crash utility to cause a segmentation fault.

You can work around the problem by loading the crash extension file only once. As a result, segmentation faults no longer occur in the described scenario.

[Bugzilla:1906482](#)

Connections fail when attaching a virtual function to virtual machine

Pensando network cards that use the **ionic** device driver silently accept VLAN tag configuration requests and attempt configuring network connections while attaching network virtual functions (**VF**) to a virtual machine (**VM**). Such network connections fail as this feature is not yet supported by the card's firmware.

Bugzilla:1930576

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.

Bugzilla:1866402

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.

Bugzilla:1793389

Using **irqpoll** causes **vmcore** generation failure

Due to an existing problem with the **nvme** driver on the 64-bit ARM architecture that run on the Amazon Web Services Graviton 1 processor, causes **vmcore** generation to fail when you provide the **irqpoll** kernel command line parameter to the first kernel. Consequently, no **vmcore** file is dumped in the **/var/crash/** directory upon a kernel crash. To work around this problem:

1. Append **irqpoll** to **KDUMP_COMMANDLINE_REMOVE** variable in the **/etc/sysconfig/kdump** file.

```
# KDUMP_COMMANDLINE_REMOVE="hugepages hugepagesz slub_debug quiet  
log_buf_len swiotlb"
```

2. Remove **irqpoll** from **KDUMP_COMMANDLINE_APPEND** variable in the **/etc/sysconfig/kdump** file.

```
# KDUMP_COMMANDLINE_APPEND="irqpoll nr_cpus=1 reset_devices  
cgroup_disable=memory udev.children-max=2 panic=10 swiotlb=noforce novmcoredd"
```

3. Restart the **kdump** service:

```
# systemctl restart kdump
```

As a result, the first kernel boots correctly and the **vmcore** file is expected to be captured upon the kernel crash.

Note that the Amazon Web Services Graviton 2 and Amazon Web Services Graviton 3 processors do not require you to manually remove the **irqpoll** parameter in the **/etc/sysconfig/kdump** file.

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.

For related information on this Known Issue, see [The **irqpoll** kernel command line parameter might cause **vmcore** generation failure](#) article.

Bugzilla:1654962

Debug kernel fails to boot in crash capture environment on RHEL 8

Due to the memory-intensive 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 as required. As a result, the debug kernel boots successfully in the crash capture environment.

Bugzilla:1659609

Allocating crash kernel memory fails at boot time

On some Ampere Altra systems, allocating the crash kernel memory during boot fails when the 32-bit region is disabled in BIOS settings. Consequently, the **kdump** service fails to start. This is caused by memory fragmentation in the region below 4 GB with no fragment being large enough to contain the crash kernel memory.

To work around this problem, enable the 32-bit memory region in BIOS as follows:

1. Open the BIOS settings on your system.
2. Open the **Chipset** menu.
3. Under **Memory Configuration**, enable the **Slave 32-bit** option.

As a result, crash kernel memory allocation within the 32-bit region succeeds and the **kdump** service works as expected.

Bugzilla:1940674

RoCE interfaces on IBM Z lose their IP settings due to an unexpected change of the network interface name

In RHEL 8.6 and earlier, the **udev** device manager assigns on the IBM Z platform unpredictable device names to RoCE interfaces that are enumerated by a unique identifier (UID). However, in RHEL 8.7 and later, **udev** assigns predictable device names with the **eno** prefix to these interfaces.

If you update from RHEL 8.6 or earlier to 8.7 or later, these UID-enumerated interfaces have new names and no longer match the device names in NetworkManager connection profiles. Consequently, these interfaces have no IP configuration after the update.

For workarounds you can apply before the update and a fix if you have already updated the system, see [RoCE interfaces on IBM Z lose their IP settings after updating to RHEL 8.7 or later](#) .

Bugzilla:2169382

The QAT manager leaves no spare device for LKCF

The Intel® QuickAssist Technology (QAT) manager (**qatmgr**) is a user space process, which by default uses all QAT devices in the system. As a consequence, there are no QAT devices left for the Linux Kernel Cryptographic Framework (LKCF). There is no need to work around this situation, as this behavior is expected and a majority of users will use acceleration from the user space.

Bugzilla:1920086

The Solarflare fails to create maximum number of virtual functions (VFs)

The Solarflare NICs fail to create a maximum number of VFs due to insufficient resources. You can check the maximum number of VFs that a PCIe device can create in the **/sys/bus/pci/devices/PCI_ID/sriov_totalvfs** file. To workaround this problem, you can either adjust the number of VFs or the VF MSI interrupt value to a lower value, either from **Solarflare Boot Manager** on startup, or using Solarflare **sfboot** utility. The default VF MSI interrupt value is **8**.

- To adjust the VF MSI interrupt value using **sfboot**:

```
# sfboot vf-msix-limit=2
```

**NOTE**

Adjusting VF MSI interrupt value affects the VF performance.

For more information about parameters to be adjusted accordingly, see the **Solarflare Server Adapter user guide**.

Bugzilla:1971506

Using `page_poison=1` can cause a kernel crash

When using `page_poison=1` as the kernel parameter on firmware with faulty EFI implementation, the operating system can cause the kernel to crash. By default, this option is disabled and it is not recommended to enable it, especially in production systems.

Bugzilla:2050411

The `iwl7260-firmware` breaks Wi-Fi on Intel Wi-Fi 6 AX200, AX210, and Lenovo ThinkPad P1 Gen 4

After updating the `iwl7260-firmware` or `iwl7260-wifi` driver to the version provided by RHEL 8.7 and later, the hardware gets into an incorrect internal state. reports its state incorrectly. Consequently, Intel Wifi 6 cards may not work and display the error message:

```
kernel: iwlwifi 0000:09:00.0: Failed to start RT ucode: -110
kernel: iwlwifi 0000:09:00.0: WRT: Collecting data: ini trigger 13 fired (delay=0ms)
kernel: iwlwifi 0000:09:00.0: Failed to run INIT ucode: -110
```

An unconfirmed work around is to power off the system and back on again. Do not reboot.

Bugzilla:2106341

Secure boot on IBM Power Systems does not support migration

Currently, on IBM Power Systems, logical partition (LPAR) does not boot after successful physical volume (PV) migration. As a result, any type of automated migration with secure boot enabled on a partition fails.

Bugzilla:2126777

`weak-modules` from `kmod` fails to work with module inter-dependencies

The `weak-modules` script provided by the `kmod` package determines which modules are kABI-compatible with installed kernels. However, while checking modules' kernel compatibility, `weak-modules` processes modules symbol dependencies from higher to lower release of the kernel for which they were built. As a consequence, modules with inter-dependencies built against different kernel releases might be interpreted as non-compatible, and therefore the `weak-modules` script fails to work in this scenario.

To work around the problem, build or put the extra modules against the latest stock kernel before you install the new kernel.

Bugzilla:2103605

kdump in Ampere Altra servers enters the OOM state

The firmware in Ampere Altra and Altra Max servers currently causes the kernel to allocate too many event, interrupt and command queues, which consumes too much memory. As a consequence, the **kdump** kernel enters the Out of memory (OOM) state.

To work around this problem, reserve extra memory for **kdump** by increasing the value of the **crashkernel=** kernel option to *640M*.

Bugzilla:2111855

Hardware certification of the real-time kernel on systems with large core-counts might require passing the skew-tick=1 boot parameter to avoid lock contentions

Large or moderate sized systems with numerous sockets and large core-counts can experience latency spikes due to lock contentions on **xtime_lock**, which is used in the timekeeping system. As a consequence, latency spikes and delays in hardware certifications might occur on multiprocessing systems. As a workaround, you can offset the timer tick per CPU to start at a different time by adding the **skew_tick=1** boot parameter.

To avoid lock conflicts, enable **skew_tick=1**:

1. Enable the **skew_tick=1** parameter with **grubby**.

```
# grubby --update-kernel=ALL --args="skew_tick=1"
```

2. Reboot for changes to take effect.
3. Verify the new settings by running the **cat /proc/cmdline** command.

Note that enabling **skew_tick=1** causes a significant increase in power consumption and, therefore, it must be enabled only if you are running latency sensitive real-time workloads.

Bugzilla:2214508

11.9. BOOT LOADER

The behavior of grubby diverges from its documentation

When you add a new kernel using the **grubby** tool and do not specify any arguments, **grubby** passes the default arguments to the new entry. This behavior occurs even without passing the **--copy-default** argument. Using **--args** and **--copy-default** options ensures those arguments are appended to the default arguments as stated in the **grubby** documentation.

However, when you add additional arguments, such as **\$tuned_params**, the **grubby** tool does not pass these arguments unless the **--copy-default** option is invoked.

In this situation, two workarounds are available:

- Either set the **root=** argument and leave **--args** empty:

```
# grubby --add-kernel /boot/my_kernel --initrd /boot/my_initrd --args "root=/dev/mapper/rhel-root" --title "entry_with_root_set"
```

- Or set the **root=** argument and the specified arguments, but not the default ones:

```
# grubby --add-kernel /boot/my_kernel --initrd /boot/my_initrd --args "root=/dev/mapper/rhel-  
root some_args and_some_more" --title "entry_with_root_set_and_other_args_too"
```

[Bugzilla:1900829](#)

11.10. FILE SYSTEMS AND STORAGE

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 **raid1**, see [Converting a mirrored LVM device to a RAID1 device](#) .

[Bugzilla:1730502](#)

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.

[Bugzilla: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.

[Bugzilla:1768536](#)

Limitations of LVM writecache

The **writocache** LVM caching method has the following limitations, which are not present in the **cache** method:

- You cannot name a **writocache** logical volume when using **pvmove** commands.
- You cannot use logical volumes with **writocache** 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 **writocache** is attached to it.

Jira:RHELPLAN-27987, [Bugzilla:1798631](#), Bugzilla:1808012

Device-mapper multipath is not supported when using NVMe/TCP driver.

The use of device-mapper multipath on top of NVMe/TCP devices can cause reduced performance and error handling. To avoid this problem, use native NVMe multipath instead of DM multipath tools. For RHEL 8, you can add the option **nvme_core.multipath=Y** to the kernel command line.

Bugzilla:2022359

The blk-availability systemd service deactivates complex device stacks

In **systemd**, the default block deactivation code does not always handle complex stacks of virtual block devices correctly. In some configurations, virtual devices might not be removed during the shutdown, which causes error messages to be logged. To work around this problem, deactivate complex block device stacks by executing the following command:

```
# systemctl enable --now blk-availability.service
```

As a result, complex virtual device stacks are correctly deactivated during shutdown and do not produce error messages.

Bugzilla:2011699

XFS quota warnings are triggered too often

Using the quota timer results in quota warnings triggering too often, which causes soft quotas to be enforced faster than they should. To work around this problem, do not use soft quotas, which will prevent triggering warnings. As a result, the amount of warning messages will not enforce soft quota limit anymore, respecting the configured timeout.

Bugzilla:2059262

11.11. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

Creating virtual Python 3.11 environments fails when using the **virtualenv** utility

The **virtualenv** utility in RHEL 8, provided by the **python3-virtualenv** package, is not compatible with Python 3.11. An attempt to create a virtual environment by using **virtualenv** will fail with the following error message:

```
$ virtualenv -p python3.11 venv3.11
Running virtualenv with interpreter /usr/bin/python3.11
ERROR: Virtual environments created by virtualenv < 20 are not compatible with Python 3.11.
```

ERROR: Use ``python3.11 -m venv`` instead.

To create Python 3.11 virtual environments, use the **python3.11 -m venv** command instead, which uses the **venv** module from the standard library.

[Bugzilla:2165702](#)

python3.11-lxml does not provide the lxml.isoschematron submodule

The **python3.11-lxml** package is distributed without the **lxml.isoschematron** submodule because it is not under an open source license. The submodule implements ISO Schematron support. As an alternative, pre-ISO-Schematron validation is available in the **lxml.etree.Schematron** class. The remaining content of the **python3.11-lxml** package is unaffected.

[Bugzilla:2157673](#)

PAM plug-in version 1.0 does not work in MariaDB

MariaDB 10.3 provides the Pluggable Authentication Modules (PAM) plug-in version 1.0. **MariaDB 10.5** provides the plug-in versions 1.0 and 2.0, version 2.0 is the default.

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.

[Bugzilla:1942330](#)

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.

Since the RHEL 8.3 update to the Apache Portable Runtime (APR) library, you can work around the problem by setting the **APR_DEEPBIND** environment variable, which enables the use of the **RTLD_DEEPBIND** dynamic linker option when loading **httpd** modules. When the **APR_DEEPBIND** environment variable is enabled, crashes no longer occur in **httpd** configurations that load conflicting libraries.

[Bugzilla:1819607](#)

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.

[Bugzilla:1803161](#)

11.12. IDENTITY MANAGEMENT

Actions required when running Samba as a print server and updating from RHEL 8.4 and earlier

With this update, the **samba** package no longer creates the **/var/spool/samba/** directory. If you use Samba as a print server and use **/var/spool/samba/** in the **[printers]** share to spool print jobs, SELinux prevents Samba users from creating files in this directory. Consequently, print jobs fail and the **auditd**

service logs a **denied** message in `/var/log/audit/audit.log`. To avoid this problem after updating your system from 8.4 and earlier:

1. Search the **[printers]** share in the `/etc/samba/smb.conf` file.
2. If the share definition contains `path = /var/spool/samba/`, update the setting and set the **path** parameter to `/var/tmp/`.
3. Restart the **smbd** service:

```
# systemctl restart smbd
```

If you newly installed Samba on RHEL 8.5 or later, no action is required. The default `/etc/samba/smb.conf` file provided by the **samba-common** package in this case already uses the `/var/tmp/` directory to spool print jobs.

Bugzilla:2009213

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.

Bugzilla:1729215

FIPS mode does not support using a shared secret to establish a cross-forest trust

Establishing a cross-forest trust using a shared secret fails in FIPS mode because NTLMSSP authentication is not FIPS-compliant. To work around this problem, authenticate with an Active Directory (AD) administrative account when establishing a trust between an IdM domain with FIPS mode enabled and an AD domain.

Bugzilla:1924707

Downgrading `authselect` after the rebase to version 1.2.2 breaks system authentication

The **authselect** package has been rebased to the latest upstream version **1.2.2**. Downgrading **authselect** is not supported and breaks system authentication for all users, including **root**.

If you downgraded the **authselect** package to **1.2.1** or earlier, perform the following steps to work around this problem:

1. At the GRUB boot screen, select **Red Hat Enterprise Linux** with the version of the kernel that you want to boot and press **e** to edit the entry.
2. Type **single** as a separate word at the end of the line that starts with **linux** and press **Ctrl+X** to start the boot process.
3. Upon booting in single-user mode, enter the root password.
4. Restore `authselect` configuration using the following command:

```
# authselect select sssd --force
```

Bugzilla:1892761

IdM to AD cross-realm TGS requests fail

The Privilege Attribute Certificate (PAC) information in IdM Kerberos tickets is now signed with AES SHA-2 HMAC encryption, which is not supported by Active Directory (AD).

Consequently, IdM to AD cross-realm TGS requests, that is, two-way trust setups, are failing with the following error:

```
Generic error (see e-text) while getting credentials for <service principal>
```

[Bugzilla:2125182](#)

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

The default keyword for enabled ciphers in the NSS does not work in conjunction with other ciphers

In Directory Server you can use the `default` keyword to refer to the default ciphers enabled in the network security services (NSS). However, if you want to enable the default ciphers and additional ones using the command line or web console, Directory Server fails to resolve the `default` keyword. As a consequence, the server enables only the additionally specified ciphers and logs an error similar to the following:

```
Security Initialization - SSL alert: Failed to set SSL cipher preference information: invalid ciphers <default,+cipher_name>: format is +cipher1,-cipher2... (Netscape Portable Runtime error 0 - no error)
```

As a workaround, specify all ciphers that are enabled by default in NSS including the ones you want to additionally enable.

[Bugzilla:1817505](#)

`pki-core-debuginfo` update from RHEL 8.6 to RHEL 8.7 or later fails

Updating the `pki-core-debuginfo` package from RHEL 8.6 to RHEL 8.7 or later fails. To work around this problem, run the following commands:

1. `yum remove pki-core-debuginfo`
2. `yum update -y`
3. `yum install pki-core-debuginfo`

4. `yum install idm-pki-symkey-debuginfo idm-pki-tools-debuginfo`

[Bugzilla:2134093](#)

Migrated IdM users might be unable to log in due to mismatching domain SIDs

If you have used the `ipa migrate-ds` script to migrate users from one IdM deployment to another, those users might have problems using IdM services because their previously existing Security Identifiers (SIDs) do not have the domain SID of the current IdM environment. For example, those users can retrieve a Kerberos ticket with the `kinit` utility, but they cannot log in. To work around this problem, see the following Knowledgebase article: [Migrated IdM users unable to log in due to mismatching domain SIDs](#).

Jira:RHELPLAN-109613

IdM in FIPS mode does not support using the NTLMSSP protocol to establish a two-way cross-forest trust

Establishing a two-way cross-forest trust between Active Directory (AD) and Identity Management (IdM) with FIPS mode enabled fails because the New Technology LAN Manager Security Support Provider (NTLMSSP) authentication is not FIPS-compliant. IdM in FIPS mode does not accept the RC4 NTLM hash that the AD domain controller uses when attempting to authenticate.

[Bugzilla:2120572](#)

IdM Vault encryption and decryption fails in FIPS mode

The OpenSSL RSA-PKCS1v15 padding encryption is blocked if FIPS mode is enabled. Consequently, Identity Management (IdM) Vaults fail to work correctly as IdM is currently using the PKCS1v15 padding for wrapping the session key with the transport certificate.

[Bugzilla:2122919](#)

Incorrect warning when setting expiration dates for a Kerberos principal

If you set a password expiration date for a Kerberos principal, the current timestamp is compared to the expiration timestamp using a 32-bit signed integer variable. If the expiration date is more than 68 years in the future, it causes an integer variable overflow resulting in the following warning message being displayed:

Warning: Your password will expire in less than one hour on [expiration date]

You can ignore this message, the password will expire correctly at the configured date and time.

[Bugzilla:2125318](#)

11.13. 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.

[Bugzilla:1668760](#)

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 or later as the host.

Bugzilla:1583445

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.

Bugzilla:1717947

11.14. GRAPHICS INFRASTRUCTURES

The radeon driver 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 system and **kdump**. After starting **kdump**, the **force_rebuild 1** line might be removed from the configuration file.

Note that in this scenario, no graphics is available during the dump process, but **kdump** works correctly.

Bugzilla: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.

Bugzilla:1812577

GUI in ESXi might crash due to low video memory

The graphical user interface (GUI) on RHEL virtual machines (VMs) in the VMware ESXi 7.0.1 hypervisor with vCenter Server 7.0.1 requires a certain amount of video memory. If you connect multiple consoles or high-resolution monitors to the VM, the GUI requires at least 16 MB of video memory. If you start the GUI with less video memory, the GUI might terminate unexpectedly.

To work around the problem, configure the hypervisor to assign at least 16 MB of video memory to the VM. As a result, the GUI on the VM no longer crashes.

If you encounter this issue, Red Hat recommends that you report it to VMware.

See also the following VMware article: [VMs with high resolution VM console may experience a crash on ESXi 7.0.1 \(83194\)](#).

Bugzilla:1910358

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.

[Bugzilla:1886147](#)

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.

[Bugzilla:1673073](#)

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 installer freezes on servers with ASPEED 2600

When you start the graphical RHEL 8.8 installer on a server with the ASPEED 2600 On System Management Chipset, the installer becomes unresponsive with a black screen. Consequently, you cannot install RHEL 8.8 on the server.

To work around the issue, add either of the following options on the kernel command line when booting the installer:

- **nomodeset**
- **drm_kms_helper.edid_firmware=edid/1024x768.bin**

As a result, the installation proceeds as expected.

Bugzilla:2189645

11.15. THE WEB CONSOLE

VNC console works incorrectly at certain resolutions

When using the Virtual Network Computing (VNC) console under certain display resolutions, you might experience a mouse offset issue or you might see only a part of the interface. Consequently, using the VNC console might not be possible. To work around this issue, you can try expanding the size of the VNC console or use the Desktop Viewer in the console tab to launch the remote viewer instead.

[Bugzilla:2030836](#)

11.16. RED HAT ENTERPRISE LINUX SYSTEM ROLES

Unable to manage localhost by using the localhost hostname in the playbook or inventory

With the inclusion of the **ansible-core 2.13** package in RHEL, if you are running Ansible on the same host you manage your nodes, you cannot do it by using the **localhost** hostname in your playbook or inventory. This happens because **ansible-core 2.13** uses the **python38** module, and many of the libraries are missing, for example, **blivet** for the **storage** role, **gobject** for the **network** role. To work around this problem, if you are already using the **localhost** hostname in your playbook or inventory, you can add a connection, by using **ansible_connection=local**, or by creating an inventory file that lists **localhost** with the **ansible_connection=local** option. With that, you are able to manage resources on **localhost**. For more details, see the article [RHEL System Roles playbooks fail when run on localhost](#) .

[Bugzilla:2041997](#)

If firewalld.service is masked, using the firewall RHEL System Role fails

If **firewalld.service** is masked on a RHEL system, the **firewall** RHEL System Role fails. To work around this problem, unmask the **firewalld.service**:

```
systemctl unmask firewalld.service
```

[Bugzilla:2123859](#)

The rhc system role fails on already registered systems when rhc_auth contains activation keys

Executing playbook files on already registered systems fails if activation keys are specified for the **rhc_auth** parameter. To work around this issue, do not specify activation keys when executing the playbook file on the already registered system.

[Bugzilla:2186908](#)

11.17. VIRTUALIZATION

Using a large number of queues might cause Windows virtual machines to fail

Windows virtual machines (VMs) might fail when the virtual Trusted Platform Module (vTPM) device is enabled and the *multi-queue virtio-net* feature is configured to use more than 250 queues.

This problem is caused by a limitation in the vTPM device. The vTPM device has a hardcoded limit on the maximum number of opened file descriptors. Since multiple file descriptors are opened for every new queue, the internal vTPM limit can be exceeded, causing the VM to fail.

To work around this problem, choose one of the following two options:

- Keep the vTPM device enabled, but use less than 250 queues.
- Disable the vTPM device to use more than 250 queues.

[Bugzilla:2020133](#)

The Milan VM CPU type is sometimes not available on AMD Milan systems

On certain AMD Milan systems, the Enhanced REP MOVSB (**erms**) and Fast Short REP MOVSB (**fstrm**) feature flags are disabled in the BIOS by default. Consequently, the **Milan** CPU type might not be available on these systems. In addition, VM live migration between Milan hosts with different feature flag settings might fail. To work around these problems, manually turn on **erms** and **fstrm** in the BIOS of your host.

[Bugzilla:2077770](#)

SMT CPU topology is not detected by VMs when using host passthrough mode on AMD EPYC

When a virtual machine (VM) boots with the CPU host passthrough mode on an AMD EPYC host, the **TOPOEXT** CPU feature flag is not present. Consequently, the VM is not able to detect a virtual CPU topology with multiple threads per core. To work around this problem, boot the VM with the EPYC CPU model instead of host passthrough.

[Bugzilla:1740002](#)

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'**.

[Bugzilla:1777138](#)

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.

[Bugzilla:1719687](#)

Virtual machines with `iommu_platform=on` fail to start on IBM POWER

RHEL 8 currently does not support the **iommu_platform=on** parameter for virtual machines (VMs) on IBM POWER system. As a consequence, starting a VM with this parameter on IBM POWER hardware results in the VM becoming unresponsive during the boot process.

[Bugzilla:1910848](#)

IBM POWER hosts now work correctly when using the `ibmvfc` driver

When running RHEL 8 on a PowerVM logical partition (LPAR), a variety of errors could previously occur due to problems with the **ibmvfc** driver. As a consequence, a kernel panic triggered on the host under certain circumstances, such as:

- Using the Live Partition Mobility (LPM) feature
- Resetting a host adapter

- Using SCSI error handling (SCSI EH) functions

With this update, the handling of **ibmvfc** has been fixed, and the described kernel panics no longer occur.

Bugzilla:1961722

Using **perf kvm record** on IBM POWER Systems can cause the VM to crash

When using a RHEL 8 host on the little-endian variant of IBM POWER hardware, using the **perf kvm record** command to collect trace event samples for a KVM virtual machine (VM) in some cases results in the VM becoming unresponsive. This situation occurs when:

- The **perf** utility is used by an unprivileged user, and the **-p** option is used to identify the VM - for example **perf kvm record -e trace_cycles -p 12345**.
- The VM was started using the **virsh** shell.

To work around this problem, use the **perf kvm** utility with the **-i** option to monitor VMs that were created using the **virsh** shell. For example:

```
# perf kvm record -e trace_imc/trace_cycles/ -p <guest pid> -i
```

Note that when using the **-i** option, child tasks do not inherit counters, and threads will therefore not be monitored.

Bugzilla:1924016

Windows Server 2016 virtual machines with Hyper-V enabled fail to boot when using certain CPU models

Currently, it is not possible to boot a virtual machine (VM) that uses Windows Server 2016 as the guest operating system, has the Hyper-V role enabled, and uses one of the following CPU models:

- EPYC-IBPB
- EPYC

To work around this problem, use the **EPYC-v3** CPU model, or manually enable the **xsaves** CPU flag for the VM.

Bugzilla:1942888

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.

Bugzilla:1741436

Using **virt-customize** sometimes causes **guestfs-firstboot** to fail

After modifying a virtual machine (VM) disk image using the **virt-customize** utility, the **guestfs-firstboot** service in some cases fails due to incorrect SELinux permissions. This causes a variety of problems during VM startup, such as failing user creation or system registration.

To avoid this problem, use the **virt-customize** command with the **--selinux-relabel** option.

[Bugzilla:1554735](#)

Deleting a forward interface from a macvtap virtual network resets all connection counts of this network

Currently, deleting a forward interface from a **macvtap** virtual network with multiple forward interfaces also resets the connection status of the other forward interfaces of the network. As a consequence, the connection information in the live network XML is incorrect. Note, however, that this does not affect the functionality of the virtual network. To work around the issue, restart the **libvirtd** service on your host.

[Bugzilla:1332758](#)

Virtual machines with SLOF fail to boot in netcat interfaces

When using a netcat (**nc**) interface to access the console of a virtual machine (VM) that is currently waiting at the Slimline Open Firmware (SLOF) prompt, the user input is ignored and VM stays unresponsive. To work around this problem, use the **nc -C** option when connecting to the VM, or use a telnet interface instead.

[Bugzilla:1974622](#)

Attaching mediated devices to virtual machines in virt-manager in some cases fails

The **virt-manager** application is currently able to detect mediated devices, but cannot recognize whether the device is active. As a consequence, attempting to attach an inactive mediated device to a running virtual machine (VM) using **virt-manager** fails. Similarly, attempting to create a new VM that uses an inactive mediated device fails with a **device not found** error.

To work around this issue, use the **virsh nodedev-start** or **mdevctl start** commands to activate the mediated device before using it in **virt-manager**.

[Bugzilla:2026985](#)

RHEL 9 virtual machines fail to boot in POWER8 compatibility mode

Currently, booting a virtual machine (VM) that runs RHEL 9 as its guest operating system fails if the VM also uses CPU configuration similar to the following:

```
<cpu mode="host-model">
  <model>power8</model>
</cpu>
```

To work around this problem, do not use POWER8 compatibility mode in RHEL 9 VMs.

In addition, note that running RHEL 9 VMs is not possible on POWER8 hosts.

[Bugzilla:2035158](#)

SUID and SGID are not cleared automatically on virtiofs

When you run the **virtiofsd** service with the **killpriv_v2** feature, your system may not automatically clear the SUID and SGID permissions after performing some file-system operations. Consequently, not clearing the permissions might cause a potential security threat. To work around this issue, disable the **killpriv_v2** feature by entering the following command:

```
# virtiofsd -o no_killpriv_v2
```

Bugzilla:1966475

Restarting the OVS service on a host might block network connectivity on its running VMs

When the Open vSwitch (OVS) service restarts or crashes on a host, virtual machines (VMs) that are running on this host cannot recover the state of the networking device. As a consequence, VMs might be completely unable to receive packets.

This problem only affects systems that use the packed virtqueue format in their **virtio** networking stack.

To work around this problem, use the **packed=off** parameter in the **virtio** networking device definition to disable packed virtqueue. With packed virtqueue disabled, the state of the networking device can, in some situations, be recovered from RAM.

Bugzilla:1792683

NFS failure during VM migration causes migration failure and source VM coredump

Currently, if the NFS service or server is shut down during virtual machine (VM) migration, the source VM's QEMU is unable to reconnect to the NFS server when it starts running again. As a result, the migration fails and a coredump is initiated on the source VM. Currently, there is no workaround available.

Bugzilla:2177957

Hotplugging a Watchdog card to a virtual machine fails

Currently, if there are no PCI slots available, adding a Watchdog card to a running virtual machine (VM) fails with the following error:

```
Failed to configure watchdog
ERROR Error attempting device hotplug: internal error: No more available PCI slots
```

To work around this problem, shut down the VM before adding the Watchdog card.

Bugzilla:2173584

11.18. RHEL IN CLOUD ENVIRONMENTS

Setting static IP in a RHEL virtual machine on a VMware host does not work

Currently, when using RHEL 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.

To work around this issue, see the [VMware knowledgebase](#).

Bugzilla:1750862

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
```


Bugzilla:1865745

The SCSI host address sometimes changes when booting a Hyper-V VM with multiple guest disks

Currently, when booting a RHEL 8 virtual machine (VM) on the Hyper-V hypervisor, the host portion of the *Host, Bus, Target, Lun* (HBTL) SCSI address in some cases changes. As a consequence, automated tasks set up with the HBTL SCSI identification or device node in the VM do not work consistently. This occurs if the VM has more than one disk or if the disks have different sizes.

To work around the problem, modify your kickstart files, using one of the following methods:

Method 1: Use persistent identifiers for SCSI devices.

You can use for example the following powershell script to determine the specific device identifiers:

```
# Output what the /dev/disk/by-id/<value> for the specified hyper-v virtual disk.
# Takes a single parameter which is the virtual disk file.
# Note: kickstart syntax works with and without the /dev/ prefix.
param (
    [Parameter(Mandatory=$true)][string]$virtualdisk
)

$what = Get-VHD -Path $virtualdisk
$part = $what.DiskIdentifier.ToLower().split('-')

$p = $part[0]
$s0 = $part[6] + $part[7] + $part[4] + $part[5] + $part[2] + $part[3] + $part[0] + $part[1]

$p = $part[1]
$s1 = $part[2] + $part[3] + $part[0] + $part[1]

[string]::format("/dev/disk/by-id/wwn-0x60022480{0}{1}{2}", $s0, $s1, $part[4])
```

You can use this script on the hyper-v host, for example as follows:

```
PS C:\Users\Public\Documents\Hyper-V\Virtual hard disks> .\by-id.ps1 .\Testing_8\disk_3_8.vhdx
/dev/disk/by-id/wwn-0x60022480e00bc367d7fd902e8bf0d3b4
PS C:\Users\Public\Documents\Hyper-V\Virtual hard disks> .\by-id.ps1 .\Testing_8\disk_3_9.vhdx
/dev/disk/by-id/wwn-0x600224807270e09717645b1890f8a9a2
```

Afterwards, the disk values can be used in the kickstart file, for example as follows:

```
part / --fstype=xfst --grow --asprimary --size=8192 --ondisk=/dev/disk/by-id/wwn-
0x600224807270e09717645b1890f8a9a2
part /home --fstype="xfs" --grow --ondisk=/dev/disk/by-id/wwn-
0x60022480e00bc367d7fd902e8bf0d3b4
```

As these values are specific for each virtual disk, the configuration needs to be done for each VM instance. It may, therefore, be useful to use the **%include** syntax to place the disk information into a separate file.

Method 2: Set up device selection by size.

A kickstart file that configures disk selection based on size must include lines similar to the following:

-

```

...

# Disk partitioning information is supplied in a file to kick start
#include /tmp/disks

...

# Partition information is created during install using the %pre section


```

%pre --interpreter /bin/bash --log /tmp/ks_pre.log

Dump whole SCSI/IDE disks out sorted from smallest to largest ouputting
just the name
disks=(`lsblk -n -o NAME -l -b -x SIZE -d -l 8,3`) || exit 1

We are assuming we have 3 disks which will be used
and we will create some variables to represent
d0=${disks[0]}
d1=${disks[1]}
d2=${disks[2]}

echo "part /home --fstype="xfs" --ondisk=$d2 --grow" >> /tmp/disks
echo "part swap --fstype="swap" --ondisk=$d0 --size=4096" >> /tmp/disks
echo "part / --fstype="xfs" --ondisk=$d1 --grow" >> /tmp/disks
echo "part /boot --fstype="xfs" --ondisk=$d1 --size=1024" >> /tmp/disks

%end

```


```

Bugzilla:1906870

RHEL instances on Azure fail to boot if provisioned by cloud-init and configured with an NFSv3 mount entry

Currently, booting a RHEL virtual machine (VM) on the Microsoft Azure cloud platform fails if the VM was provisioned by the **cloud-init** tool and the guest operating system of the VM has an NFSv3 mount entry in the **/etc/fstab** file.

Bugzilla:2081114

11.19. SUPPORTABILITY

The **getattachment** command fails to download multiple attachments at once

The **redhat-support-tool** command offers the **getattachment** subcommand for downloading attachments. However, **getattachment** is currently only able to download a single attachment and fails to download multiple attachments.

As a workaround, you can download multiple attachments one by one by passing the case number and UUID for each attachment in the **getattachment** subcommand.

Bugzilla:2064575

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.

[Bugzilla:1802026](#)

Timeout when running `sos report` on IBM Power Systems, Little Endian

When running the **sos report** command on IBM Power Systems, Little Endian with hundreds or thousands of CPUs, the processor plugin reaches its default timeout of 300 seconds when collecting huge content of the `/sys/devices/system/cpu` directory. As a workaround, increase the plugin's timeout accordingly:

- For one-time setting, run:

```
# sos report -k processor.timeout=1800
```

- For a permanent change, edit the **[plugin_options]** section of the `/etc/sos/sos.conf` file:

```
[plugin_options]
# Specify any plugin options and their values here. These options take the form
# plugin_name.option_name = value
#rpm.rpmva = off
processor.timeout = 1800
```

The example value is set to 1800. The particular timeout value highly depends on a specific system. To set the plugin's timeout appropriately, you can first estimate the time needed to collect the one plugin with no timeout by running the following command:

```
# time sos report -o processor -k processor.timeout=0 --batch --build
```

[Bugzilla:2011413](#)

11.20. CONTAINERS

Running `systemd` within an older container image does not work

Running `systemd` within an older container image, for example, **centos:7**, does not work:

```
$ podman run --rm -ti centos:7 /usr/lib/systemd/systemd
Storing signatures
Failed to mount cgroup at /sys/fs/cgroup/systemd: Operation not permitted
[!!!!!!] Failed to mount API filesystems, freezing.
```

To work around this problem, use the following commands:

```
# mkdir /sys/fs/cgroup/systemd
# mount none -t cgroup -o none,name=systemd /sys/fs/cgroup/systemd
# podman run --runtime /usr/bin/crun --annotation=run.oci.systemd.force_cgroup_v1=/sys/fs/cgroup -
-rm -ti centos:7 /usr/lib/systemd/systemd
```

[Jira:RHELPLAN-96940](#)

CHAPTER 12. INTERNATIONALIZATION

12.1. 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

12.2. 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 tickets are listed in this document for reference. The links lead to the release notes in this document that describe the tickets.

Component	Tickets
389-ds-base	Bugzilla:2136610 , Bugzilla:2096795 , Bugzilla:2142639 , Bugzilla:2130276 , Bugzilla:1817505
NetworkManager	Bugzilla:2089707 , Bugzilla:2134907 , Bugzilla:2132754
SLOF	Bugzilla:1910848
accel-config	Bugzilla:1843266
anaconda	Bugzilla:1913035 , Bugzilla:2014103 , Bugzilla:1991516 , Bugzilla:2094977 , Bugzilla:2050140 , Bugzilla:1914955 , Bugzilla:1929105 , Bugzilla:2126506
ansible-collection-microsoft-sql	Bugzilla:2144820 , Bugzilla:2144821 , Bugzilla:2144852 , Bugzilla:2153428 , Bugzilla:2163696 , Bugzilla:2153427
ansible-freeipa	Bugzilla:2127912
apr	Bugzilla:1819607
authselect	Bugzilla:1892761
bacula	Bugzilla:2089399
brltty	Bugzilla:2008197
certmonger	Bugzilla:2150025
clevis	Bugzilla:2159440 , Bugzilla:2159736
cloud-init	Bugzilla:1750862
cockpit	Bugzilla:2212371 , Bugzilla:1666722
cockpit-appstream	Bugzilla:2030836
cockpit-machines	Bugzilla:2173584
conntrack-tools	Bugzilla:2126736
coreutils	Bugzilla:2030661

Component	Tickets
corosync-qdevice	Bugzilla:1784200
crash	Bugzilla:1906482
crash-ptdump-command	Bugzilla:1838927
createrepo_c	Bugzilla:1973588
crypto-policies	Bugzilla:1921646 , Bugzilla:2071981 , Bugzilla:1919155 , Bugzilla:1660839
device-mapper-multipath	Bugzilla:2022359 , Bugzilla:2011699
distribution	Bugzilla:1657927
dnf	Bugzilla:2054235 , Bugzilla:2047251 , Bugzilla:2016070 , Bugzilla:1986657
dnf-plugins-core	Bugzilla:2139324
edk2	Bugzilla:1741615 , Bugzilla:1935497
fapolicyd	Bugzilla:2165645 , Bugzilla:2054741
fence-agents	Bugzilla:1775847
firewalld	Bugzilla:1871860
gcc	Bugzilla:2110582
gdb	Bugzilla:1853140
git	Bugzilla:2139378
git-lfs	Bugzilla:2139382
glassfish-jaxb	Bugzilla:2055539
glibc	Bugzilla:1871383 , Bugzilla:1159809
gnome-session	Bugzilla:2070976
gnome-shell-extensions	Bugzilla:2033572 , Bugzilla:2138109 , Bugzilla:1717947
gnome-software	Bugzilla:1668760

Component	Tickets
gnutls	Bugzilla:1628553
golang	Bugzilla:2174430 , Bugzilla:2132767 , Bugzilla:2132694 , Bugzilla:2132419
grub2	Bugzilla:1583445
grubby	Bugzilla:1900829
initscripts	Bugzilla:1875485
ipa	Bugzilla:2075452 , Bugzilla:1924707 , Bugzilla:2120572 , Bugzilla:2122919 , Bugzilla:1664719 , Bugzilla:1664718 , Bugzilla:2101770
ipmitool	Bugzilla:1873614
kernel	Bugzilla:2107595 , Bugzilla:1660908 , Bugzilla:1664379 , Bugzilla:2136107 , Bugzilla:2127136 , Bugzilla:2143849 , Bugzilla:1905243 , Bugzilla:2009705 , Bugzilla:2103946 , Bugzilla:2087262 , Bugzilla:2151854 , Bugzilla:2134931 , Bugzilla:2069047 , Bugzilla:2135417 , Bugzilla:1868526 , Bugzilla:1694705 , Bugzilla:1730502 , Bugzilla:1609288 , Bugzilla:1602962 , Bugzilla:1865745 , Bugzilla:1906870 , Bugzilla:1924016 , Bugzilla:1942888 , Bugzilla:1812577 , Bugzilla:1910358 , Bugzilla:1930576 , Bugzilla:1793389 , Bugzilla:1654962 , Bugzilla:1940674 , Bugzilla:2169382 , Bugzilla:1920086 , Bugzilla:1971506 , Bugzilla:2059262 , Bugzilla:2050411 , Bugzilla:2106341 , Bugzilla:2127028 , Bugzilla:2130159 , Bugzilla:2189645 , Bugzilla:1605216 , Bugzilla:1519039 , Bugzilla:1627455 , Bugzilla:1501618 , Bugzilla:1633143 , Bugzilla:1814836 , Bugzilla:1839311 , Bugzilla:1570255 , Bugzilla:1696451 , Bugzilla:1348508 , Bugzilla:1837187 , Bugzilla:1660337 , Bugzilla:2041686 , Bugzilla:1836977 , Bugzilla:1878207 , Bugzilla:1665295 , Bugzilla:1871863 , Bugzilla:1569610 , Bugzilla:1794513
kexec-tools	Bugzilla:2111855
kmod	Bugzilla:2103605
kmod-kvdo	Bugzilla:2119819 , Bugzilla:2109047
krb5	Bugzilla:2125182 , Bugzilla:2125318 , Bugzilla:1877991
libdnf	Bugzilla:2124483
libffi	Bugzilla:2014228
libgnome-keyring	Bugzilla:1607766
libguestfs	Bugzilla:1554735

Component	Tickets
libreswan	Bugzilla:2128672 , Bugzilla:2176248 , Bugzilla:1989050
libselinux-python-2.8-module	Bugzilla:1666328
libsoup	Bugzilla:1938011
libvirt	Bugzilla:1664592 , Bugzilla:1332758 , Bugzilla:1528684
llvm-toolset	Bugzilla:2118568
lvm2	Bugzilla:1496229 , Bugzilla:1768536
mariadb	Bugzilla:1942330
mesa	Bugzilla:1886147
mod_security	Bugzilla:2143207
nfs-utils	Bugzilla:2081114 , Bugzilla:1592011
nginx	Bugzilla:2112345
nispor	Bugzilla:2153166
nodejs	Bugzilla:2178087
nss	Bugzilla:1817533 , Bugzilla:1645153
nss_nis	Bugzilla:1803161
openblas	Bugzilla:2115722
opencryotoki	Bugzilla:2110315
opencv	Bugzilla:1886310
openmpi	Bugzilla:1866402
opensc	Bugzilla:2176973 , Bugzilla:1947025 , Bugzilla:2097048
openscap	Bugzilla:2159290 , Bugzilla:2161499
openssh	Bugzilla:2044354

Component	Tickets
openssl	Bugzilla:1810911
oscap-anaconda-addon	Bugzilla:2075508 , Bugzilla:1843932 , Bugzilla:1665082 , Bugzilla:2165948
pacemaker	Bugzilla:2133497 , Bugzilla:2121852 , Bugzilla:2122806
pam	Bugzilla:2068461
pcs	Bugzilla:2132582 , Bugzilla:1816852 , Bugzilla:2112263 , Bugzilla:2112267 , Bugzilla:1918527 , Bugzilla:1619620 , Bugzilla:1851335
pki-core	Bugzilla:1729215 , Bugzilla:2134093 , Bugzilla:1628987
podman	Jira:RHELPLAN-136601 , Jira:RHELPLAN-136608 , Bugzilla:2119200 , Jira:RHELPLAN-136610
postfix	Bugzilla:1711885
postgresql	Bugzilla:2128241
powertop	Bugzilla:2040070
pykickstart	Bugzilla:1637872
python3.11	Bugzilla:2137139
python3.11-lxml	Bugzilla:2157673
python36-3.6-module	Bugzilla:2165702
qemu-kvm	Bugzilla:2117149 , Bugzilla:2020133 , Bugzilla:1740002 , Bugzilla:1719687 , Bugzilla:1966475 , Bugzilla:1792683 , Bugzilla:2177957 , Bugzilla:1651994
rear	Bugzilla:2130206 , Bugzilla:2172605 , Bugzilla:2131946 , Bugzilla:1925531 , Bugzilla:2083301
redhat-support-tool	Bugzilla:2064575 , Bugzilla:1802026
restore	Bugzilla:1997366

Component	Tickets
rhel-system-roles	Bugzilla:2119600, Bugzilla:2130019, Bugzilla:2143814, Bugzilla:2079009, Bugzilla:2130332, Bugzilla:2130345, Bugzilla:2133532, Bugzilla:2133931, Bugzilla:2134201, Bugzilla:2133856, Bugzilla:2143458, Bugzilla:2137667, Bugzilla:2143385, Bugzilla:2144876, Bugzilla:2144877, Bugzilla:2130362, Bugzilla:2129620, Bugzilla:2165176, Bugzilla:2149683, Bugzilla:2126960, Bugzilla:2127497, Bugzilla:2153081, Bugzilla:2167941, Bugzilla:2153080, Bugzilla:2168733, Bugzilla:2162782, Bugzilla:2123859, Bugzilla:2186908, Bugzilla:2021685, Bugzilla:2006081
rpm	Bugzilla:2129345, Bugzilla:2110787, Bugzilla:1688849
rsync	Bugzilla:2139118
rsyslog	Bugzilla:2124934, Bugzilla:2070496, Bugzilla:2157658, Bugzilla:1679512, Jira:RHELPLAN-10431
rt-tests	Bugzilla:2122374
rteval	Bugzilla:2082260
rtla	Bugzilla:2075203
rust-toolset	Bugzilla:2123899
s390utils	Bugzilla:2043833
samba	Bugzilla:2132051, Bugzilla:2009213, Jira:RHELPLAN-13195, Jira:RHELDPCS-16612
scap-security-guide	Bugzilla:2072444, Bugzilla:2152658, Bugzilla:2156192, Bugzilla:2158404, Bugzilla:2119356, Bugzilla:2122322, Bugzilla:2115343, Bugzilla:2152208, Bugzilla:2099394, Bugzilla:2151553, Bugzilla:2162803, Bugzilla:2028428, Bugzilla:2118758, Bugzilla:2167373
selinux-policy	Bugzilla:1972230, Bugzilla:2088441, Bugzilla:2154242, Bugzilla:2134125, Bugzilla:2090711, Bugzilla:2101341, Bugzilla:2121709, Bugzilla:2122838, Bugzilla:2124388, Bugzilla:2125008, Bugzilla:2143696, Bugzilla:2148561, Bugzilla:1461914
sos	Bugzilla:2164987, Bugzilla:2134906, Bugzilla:2011413
spice	Bugzilla:1849563
sssd	Bugzilla:2144519, Bugzilla:2087247, Bugzilla:2065692, Bugzilla:2056483, Bugzilla:1947671
subscription-manager	Bugzilla:2170082

Component	Tickets
swig	Bugzilla:2139076
synce4l	Bugzilla:2019751
tang	Bugzilla:2188743
texlive	Bugzilla:2150727
tomcat	Bugzilla:2160455
tuna	Bugzilla:2121518
tuned	Bugzilla:2133814 , Bugzilla:2113900
tzdata	Bugzilla:2154109
udica	Bugzilla:1763210
usbguard	Bugzilla:2159409 , Bugzilla:2159411 , Bugzilla:2159413
vdo	Bugzilla:1949163
virt-manager	Bugzilla:2026985
wayland	Bugzilla:1673073
weldr-client	Bugzilla:2033192
wsmancli	Bugzilla:2105316
xdp-tools	Bugzilla:2160069
xorg-x11-server	Bugzilla:1698565

Component	Tickets
other	<p>Bugzilla:2177769, Jira:RHELPLAN-139125, Jira:RHELPLAN-137505, Jira:RHELPLAN-139430, Jira:RHELPLAN-137416, Jira:RHELPLAN-137411, Jira:RHELPLAN-137406, Jira:RHELPLAN-137403, Jira:RHELPLAN-139448, Jira:RHELPLAN-151481, Jira:RHELPLAN-150266, Jira:RHELPLAN-151121, Jira:RHELPLAN-149091, Jira:RHELPLAN-139424, Jira:RHELPLAN-136489, Bugzilla:2183445, Jira:RHELPLAN-59528, Jira:RHELPLAN-148303, Bugzilla:2025814, Bugzilla:2077770, Bugzilla:1777138, Bugzilla:1640697, Bugzilla:1697896, Bugzilla:1961722, Bugzilla:1659609, Bugzilla:1687900, Bugzilla:1757877, Bugzilla:1741436, Jira:RHELPLAN-27987, Jira:RHELPLAN-34199, Jira:RHELPLAN-57914, Jira:RHELPLAN-96940, Bugzilla:1974622, Bugzilla:2028361, Bugzilla:2041997, Bugzilla:2035158, Jira:RHELPLAN-109613, Bugzilla:2126777, Bugzilla:1690207, Bugzilla:1559616, Bugzilla:1889737, Bugzilla:1906489, Bugzilla:1769727, Jira:RHELPLAN-27394, Jira:RHELPLAN-27737, Jira:RHELPLAN-148394, Bugzilla:1642765, Bugzilla:1646541, Bugzilla:1647725, Bugzilla:1932222, Bugzilla:1686057, Bugzilla:1748980, Jira:RHELPLAN-71200, Jira:RHELPLAN-45858, Bugzilla:1871025, Bugzilla:1871953, Bugzilla:1874892, Bugzilla:1916296, Jira:RHELPLAN-100400, Bugzilla:1926114, Bugzilla:1904251, Bugzilla:2011208, Jira:RHELPLAN-59825, Bugzilla:1920624, Jira:RHELPLAN-70700, Bugzilla:1929173, Jira:RHELPLAN-85066, Jira:RHELPLAN-98983, Bugzilla:2009113, Bugzilla:1958250, Bugzilla:2038929, Bugzilla:2006665, Bugzilla:2029338, Bugzilla:2061288, Bugzilla:2060759, Bugzilla:2055826, Bugzilla:2059626, Jira:RHELPLAN-133171, Bugzilla:2142499, Jira:RHELPLAN-145958, Jira:RHELPLAN-146398, Jira:RHELPLAN-153267</p>

APPENDIX B. REVISION HISTORY

0.2-2

Fri May 10 2024, Brian Angelica (bangelic@redhat.com)

- Updated Tech Preview in [BZ#1690207](#).

0.2-1

Thu May 9 2024, Gabriela Fialova (gfialova@redhat.com)

- Updated a known issue [BZ#1730502](#) (Storage).

0.1-10

Thu Apr 25 2024, Gabriela Fialová (gfialova@redhat.com)

- Added an enhancement [BZ#2136610](#) (Identity Management).

0.1-9

Mon Mar 04 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added a bug fix [Jira:SSSD-6096](#) (Identity Management).

0.1-8

Thu Feb 29 2024, Gabriela Fialová (gfialova@redhat.com)

- Added a deprecated functionality [Jira:RHELDOCS-17641](#) (Networking).

0.1-7

Tue Feb 13 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added a deprecated functionality [Jira:RHELDOCS-17573](#) (Identity Management).

0.1-6

Fri Feb 2 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added a known issue [BZ#1834716](#) (Security).
- Updated text for [BZ#2183445](#) (Kernel).

0.1-5

Thu Dec 7 2023, Lucie Vařáková (lvarakova@redhat.com)

- Added a new feature [BZ#2043852](#) (Kernel).

0.1-4

Fri Nov 10 2023, Gabriela Fialová (gfialova@redhat.com)

- Updated the module on Providing Feedback on RHEL Documentation.

0.1-3

Tue Oct 17 2023, Gabriela Fialová (gfialova@redhat.com)

- Added a deprecated functionality [Jira:RHELDOCS-16755](#) (Identity Management).

- Update doc text of DF JIRA-RHELDPCS-16755 (Containers).

0.1-2

Fri Oct 13 2023, Gabriela Fialová (gfialova@redhat.com)

- Added a Tech Preview [JIRA:RHELDPCS-16861](#) (Containers).

0.1-1

October 9 2023, Lucie Vařáková (lvarakova@redhat.com)

- Updated a known issue [BZ#2169382](#) (kernel).

0.1-0

September 8 2023, Lucie Vařáková (lvarakova@redhat.com)

- Added a deprecated functionality release note [JIRA:RHELDPCS-16612](#) (Samba).
- Updated the [Providing feedback on Red Hat documentation](#) section.

0.0-9

August 24 2023, Lucie Vařáková (lvarakova@redhat.com)

- Added a known issue [BZ#2214508](#) (Kernel).

0.0-8

August 4 2023, Lenka Špačková (lspackova@redhat.com)

- Fixed section for [BZ#2225332](#).

0.0-7

August 3 2023, Lucie Vařáková (lvarakova@redhat.com)

- Added deprecated functionality [Jira:RHELPLAN-139456](#) (Identity Management).

0.0-6

August 1 2023, Lenka Špačková (lspackova@redhat.com)

- Added deprecated functionality [BZ#2225332](#).
- Improved abstract.

0.0-5

July 31 2023, Mirek Jahoda (mjahoda@redhat.com)

- The known issue [BZ#2203361](#) changed to a bug fix [BZ#2212371](#).

0.0-4

July 13 2023, Lucie Vařáková (lvarakova@redhat.com)

- Added a Technology Preview [BZ#1570255](#) (Networking).
- Updated the [In-place upgrade and OS conversion](#) section.

0.0-3

June 27 2023, Lucie Vařáková (lvarakova@redhat.com)

- Added an enhancement [BZ#2087247](#) (Identity Management).
- Moved [BZ#2176248](#) to Bug Fixes (Security).
- Added a known issue [BZ#2176973](#) (Security).
- Updated Technology Preview [BZ#1769727](#) (Kernel).

0.0-2

Jun 6 2023, Lucie Vařáková (lvarakova@redhat.com)

- Added a known issue [BZ#2177957](#) (Virtualization).
- Other small updates.

0.0-1

May 17 2023, Lucie Vařáková (lvarakova@redhat.com)

- Release of the Red Hat Enterprise Linux 8.8 Release Notes.

0.0-0

Mar 29 2023, Lucie Vařáková (lvarakova@redhat.com)

- Release of the Red Hat Enterprise Linux 8.8 Beta Release Notes.