



Red Hat Enterprise Linux 9

9.2 Release Notes

Release Notes for Red Hat Enterprise Linux 9.2

Red Hat Enterprise Linux 9 9.2 Release Notes

Release Notes for Red Hat Enterprise Linux 9.2

Legal Notice

Copyright © 2023 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 9.2 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 9.2	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 system toolchain	8
Updated performance tools and debuggers	8
Updated performance monitoring tools	8
Updated compiler toolsets	8
Java implementations in RHEL 9	8
The web console	9
Containers	9
1.2. IN-PLACE UPGRADE	9
In-place upgrade from RHEL 8 to RHEL 9	9
In-place upgrade from RHEL 7 to RHEL 9	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 9	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	16
4.3. SOFTWARE MANAGEMENT	17
4.4. SHELLS AND COMMAND-LINE TOOLS	18
4.5. INFRASTRUCTURE SERVICES	19
4.6. SECURITY	22
4.7. NETWORKING	27
4.8. KERNEL	32
4.9. FILE SYSTEMS AND STORAGE	37
4.10. HIGH AVAILABILITY AND CLUSTERS	39
4.11. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	39
4.12. COMPILERS AND DEVELOPMENT TOOLS	44
4.13. IDENTITY MANAGEMENT	50
4.14. DESKTOP	58
4.15. THE WEB CONSOLE	58
4.16. RED HAT ENTERPRISE LINUX SYSTEM ROLES	59
4.17. VIRTUALIZATION	64
4.18. SUPPORTABILITY	65
4.19. CONTAINERS	65

CHAPTER 5. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS	69
New kernel parameters	69
Updated kernel parameters	71
New sysctl parameters	78
Changed sysctl parameters	80
CHAPTER 6. DEVICE DRIVERS	81
6.1. NEW DRIVERS	81
Network drivers	81
Graphics drivers and miscellaneous drivers	82
6.2. UPDATED DRIVERS	82
Storage driver updates	82
CHAPTER 7. AVAILABLE BPF FEATURES	83
CHAPTER 8. BUG FIXES	102
8.1. INSTALLER AND IMAGE CREATION	102
8.2. SUBSCRIPTION MANAGEMENT	103
8.3. SOFTWARE MANAGEMENT	103
8.4. SHELLS AND COMMAND-LINE TOOLS	104
8.5. SECURITY	105
8.6. NETWORKING	108
8.7. KERNEL	109
8.8. BOOT LOADER	109
8.9. FILE SYSTEMS AND STORAGE	109
8.10. HIGH AVAILABILITY AND CLUSTERS	110
8.11. COMPILERS AND DEVELOPMENT TOOLS	111
8.12. IDENTITY MANAGEMENT	112
8.13. THE WEB CONSOLE	113
8.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES	113
8.15. VIRTUALIZATION	115
CHAPTER 9. TECHNOLOGY PREVIEWS	117
9.1. INSTALLER AND IMAGE CREATION	117
9.2. SHELLS AND COMMAND-LINE TOOLS	117
9.3. INFRASTRUCTURE SERVICES	117
9.4. SECURITY	117
9.5. NETWORKING	118
9.6. KERNEL	118
9.7. FILE SYSTEMS AND STORAGE	119
9.8. COMPILERS AND DEVELOPMENT TOOLS	120
9.9. IDENTITY MANAGEMENT	121
9.10. DESKTOP	122
9.11. GRAPHICS INFRASTRUCTURES	123
9.12. THE WEB CONSOLE	123
9.13. VIRTUALIZATION	123
9.14. RHEL IN CLOUD ENVIRONMENTS	125
9.15. CONTAINERS	125
CHAPTER 10. DEPRECATED FUNCTIONALITY	126
10.1. INSTALLER AND IMAGE CREATION	126
10.2. SUBSCRIPTION MANAGEMENT	127
10.3. SHELLS AND COMMAND-LINE TOOLS	127
10.4. SECURITY	127

10.5. NETWORKING	129
10.6. KERNEL	130
10.7. FILE SYSTEMS AND STORAGE	130
10.8. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	131
10.9. COMPILERS AND DEVELOPMENT TOOLS	131
10.10. IDENTITY MANAGEMENT	131
10.11. DESKTOP	132
10.12. GRAPHICS INFRASTRUCTURES	133
10.13. RED HAT ENTERPRISE LINUX SYSTEM ROLES	134
10.14. VIRTUALIZATION	134
10.15. CONTAINERS	135
10.16. DEPRECATED PACKAGES	136
CHAPTER 11. KNOWN ISSUES	138
11.1. INSTALLER AND IMAGE CREATION	138
11.2. SOFTWARE MANAGEMENT	141
11.3. SHELLS AND COMMAND-LINE TOOLS	142
11.4. INFRASTRUCTURE SERVICES	142
11.5. SECURITY	144
11.6. NETWORKING	148
11.7. KERNEL	149
11.8. BOOT LOADER	153
11.9. FILE SYSTEMS AND STORAGE	154
11.10. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	155
11.11. COMPILERS AND DEVELOPMENT TOOLS	155
11.12. IDENTITY MANAGEMENT	156
11.13. DESKTOP	162
11.14. GRAPHICS INFRASTRUCTURES	162
11.15. THE WEB CONSOLE	163
11.16. RED HAT ENTERPRISE LINUX SYSTEM ROLES	164
11.17. VIRTUALIZATION	164
11.18. RHEL IN CLOUD ENVIRONMENTS	168
11.19. SUPPORTABILITY	169
11.20. CONTAINERS	170
APPENDIX A. LIST OF TICKETS BY COMPONENT	171
APPENDIX B. REVISION HISTORY	179

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 comments on specific passages

1. View the documentation in the **Multi-page HTML** format and ensure that you see the **Feedback** button in the upper right corner after the page fully loads.
2. Use your cursor to highlight the part of the text that you want to comment on.
3. Click the **Add Feedback** button that appears near the highlighted text.
4. Add your feedback and click **Submit**.

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 9.2

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.
- Creating customized files and directories in the **/etc** directory is now supported.
- 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

Key highlights for RHEL for Edge:

- Specifying a user in a blueprint for **simplified-installer** images is now supported.
- The Ignition provisioning utility is now supported in RHEL for Edge Simplified images.
- Simplified Installer images can now be composed without the FDO customization section in the blueprint.

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

Security

Key security-related highlights:

- The **OpenSSL** secure communications library was rebased to version 3.0.7.
- **SELinux user-space** packages were updated to version 3.5.
- **Keylime** was rebased to version 6.5.2
- **OpenSCAP** was rebased to version 1.3.7.
- **SCAP Security Guide** was rebased to version 0.1.66.
- A new rule for idle session termination was added to the SCAP Security Guide.
- **Clevis** now accepts external tokens.
- **Rsyslog** TLS-encrypted logging now supports multiple CA files.
- Rsyslog privileges are limited to minimize security exposure.
- The **fapolicyd** framework now provides filtering of the RPM database.

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 system toolchain

The following system toolchain components have been updated in RHEL 9.2:

- **GCC 11.3.1**
- **glibc 2.34**
- **binutils 2.35.2**

Updated performance tools and debuggers

The following performance tools and debuggers have been updated in RHEL 9.2:

- **GDB 10.2**
- **Valgrind 3.19**
- **SystemTap 4.8**
- **Dyninst 12.1.0**
- **elfutils 0.188**

Updated performance monitoring tools

The following performance monitoring tools have been updated in RHEL 9.2:

- **PCP 6.0.1**
- **Grafana 9.0.9**

Updated compiler toolsets

The following compiler toolsets have been updated in RHEL 9.2:

- **GCC Toolset 12**
- **LLVM Toolset 15.0.7**
- **Rust Toolset 1.66**
- **Go Toolset 1.19.6**

For detailed changes, see [New features - Compilers and development tools](#).

Java implementations in RHEL 9

The RHEL 9 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 and RHEL 9.2 and later releases. With 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 CNI network stack has been deprecated.

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

1.2. IN-PLACE UPGRADE

In-place upgrade from RHEL 8 to RHEL 9

The supported in-place upgrade paths currently are:

- From RHEL 8.6 to RHEL 9.0 and RHEL 8.8 to RHEL 9.2 on the following architectures:

- 64-bit Intel
 - 64-bit AMD
 - 64-bit ARM
 - IBM POWER 9 (little endian)
 - IBM Z architectures, excluding z13
- From RHEL 8.6 to RHEL 9.0 and RHEL 8.8 to RHEL 9.2 on systems with SAP HANA

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 8 to RHEL 9](#) .

If you are upgrading to RHEL 9.2 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 8 to RHEL 9](#) .

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](#).
- With the release of RHEL 9.2, multiple upgrade paths are now available for the in-place upgrade from RHEL 8 to RHEL 9. For the current release, it is possible to perform an in-place upgrade from either RHEL 8.8 to RHEL 9.2, or RHEL 8.6 to RHEL 9.0. Note that the available upgrade paths differ between standard RHEL systems and RHEL systems with SAP HANA.
- The latest release of the **leapp-upgrade-el8toel9** package now contains all required leapp data files. Customers no longer need to manually download these data files.
- In-place upgrades of RHEL 8.8 systems in FIPS mode are now supported.
- 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 7 to RHEL 9

It is not possible to perform an in-place upgrade directly from RHEL 7 to RHEL 9. However, you can perform an in-place upgrade from RHEL 7 to RHEL 8 and then perform a second in-place upgrade to RHEL 9. For more information, see [Upgrading from RHEL 7 to RHEL 8](#) .

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](#)
- [Kickstart Generator](#)
- [Red Hat Product Certificates](#)
- [Red Hat CVE Checker](#)
- [Kernel Oops Analyzer](#)
- [Red Hat Code Browser](#)
- [VNC Configurator](#)
- [Red Hat OpenShift Container Platform Update Graph](#)
- [Red Hat Satellite Upgrade Helper](#)
- [JVM Options Configuration Tool](#)
- [Load Balancer Configuration Tool](#)
- [Red Hat OpenShift Data Foundation Supportability and Interoperability Checker](#)
- [Ansible Automation Platform Upgrade Assistant](#)
- [Ceph Placement Groups \(PGs\) per Pool Calculator](#)

1.4. ADDITIONAL RESOURCES

Capabilities and limits of Red Hat Enterprise Linux 9 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 9, including licenses and application compatibility levels.

Application compatibility levels are explained in the [Red Hat Enterprise Linux 9: Application Compatibility Guide](#) document.

Major **differences between RHEL 8 and RHEL 9**, including removed functionality, are documented in [Considerations in adopting RHEL 9](#) .

Instructions on how to perform an **in-place upgrade from RHEL 8 to RHEL 9** are provided by the document [Upgrading from RHEL 8 to RHEL 9](#) .

The **Red Hat Insights** service, which enables you to proactively identify, examine, and resolve known technical issues, is 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 9.2 is distributed with the kernel version 5.14.0-284.11.1, which provides support for the following architectures at the minimum required version:

- AMD and Intel 64-bit architectures (x86-64-v2)
- The 64-bit ARM architecture (ARMv8.0-A)
- IBM Power Systems, Little Endian (POWER9)
- 64-bit IBM Z (z14)

Make sure you purchase the appropriate subscription for each architecture. For more information, see [Get Started with Red Hat Enterprise Linux - additional architectures](#) .

CHAPTER 3. DISTRIBUTION OF CONTENT IN RHEL 9

3.1. INSTALLATION

Red Hat Enterprise Linux 9 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:

- **Installation ISO:** A full installation image that contains the BaseOS and AppStream repositories and allows you to complete the installation without additional repositories. On the [Product Downloads](#) page, the **Installation ISO** is referred to as **Binary DVD**.



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 Installation ISO image. You can also register to Red Hat CDN or Satellite during the installation to use the latest BaseOS and AppStream content from Red Hat CDN or Satellite.

See the [Performing a standard RHEL 9 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 9 installation](#) document.

3.2. REPOSITORIES

Red Hat Enterprise Linux 9 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 more information, see the [Scope of Coverage Details](#) document.

Content in the AppStream repository includes additional user-space applications, runtime languages, and databases in support of the varied workloads and use cases.

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 9 repositories and the packages they provide, see the [Package manifest](#).


3.3. APPLICATION STREAMS

Multiple versions of user-space components are delivered as Application Streams and updated more frequently than the core operating system packages. This provides greater flexibility to customize RHEL without impacting the underlying stability of the platform or specific deployments.

Application Streams are available in the familiar RPM format, as an extension to the RPM format called modules, as Software Collections, or as Flatpaks.

Each Application Stream component has a given life cycle, either the same as RHEL 9 or shorter. For RHEL life cycle information, see [Red Hat Enterprise Linux Life Cycle](#).

RHEL 9 improves the Application Streams experience by providing initial Application Stream versions that can be installed as RPM packages using the traditional **dnf install** command.



NOTE

Certain initial Application Streams in the RPM format have a shorter life cycle than Red Hat Enterprise Linux 9.

Some additional Application Stream versions will be distributed as modules with a shorter life cycle in future minor RHEL 9 releases. 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.

Always determine what version of an Application Stream you want to install and make sure to review the [Red Hat Enterprise Linux Application Stream Lifecycle](#) first.

Content that needs rapid updating, such as alternate compilers and container tools, is available in rolling streams that will not provide alternative versions in parallel. Rolling streams may be packaged as RPMs or modules.

For information about Application Streams available in RHEL 9 and their application compatibility level, see the [Package manifest](#). Application compatibility levels are explained in the [Red Hat Enterprise Linux 9: Application Compatibility Guide](#) document.

3.4. PACKAGE MANAGEMENT WITH YUM/DNF

In Red Hat Enterprise Linux 9, software installation is ensured by **DNF**. Red Hat continues to support the usage of the **yum** term for consistency with previous major versions of RHEL. If you type **dnf** instead of **yum**, the command works as expected because both are aliases for compatibility.

Although RHEL 8 and RHEL 9 are based on **DNF**, they are compatible with **YUM** used in RHEL 7.

For more information, see [Managing software with the DNF tool](#).

CHAPTER 4. NEW FEATURES

This part describes new features and major enhancements introduced in Red Hat Enterprise Linux 9.2.

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

Ability to create customized files and directories in the `/etc` directory

With this enhancement, two new blueprint customizations are available. The `[[customizations.files]]` and the `[[customizations.directories]]` blueprint customizations enable you to create customized files and directories in the `/etc` directory of your image. Currently, you can use these customization only in the `/etc` directory.

The `[[customizations.directories]]` enables you to:

- Create new directories

- Set user and group ownership for the directory
- Set the mode permission in the octal format

With the **[[customizations.files]]** blueprint customizations you can:

- Create new files under the parent / directory
- Modifying existing files - this overrides the existing content
- Set user and group ownership for the file you are creating
- Set the mode permission in the octal format



NOTE

The new blueprint customizations are supported by all the image types, such as **edge-container**, **edge-commit**, among others. The customizations not supported in the blueprints used to create Installer images, such as **edge-raw-image**, **edge-installer**, and **edge-simplified-installer**.

Jira:RHELPLAN-147428

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

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

Minimal RHEL installation now installs only the **s390utils-core** package

In RHEL 8.4 and later, the **s390utils-base** package is split into an **s390utils-core** package and an auxiliary **s390utils-base** package. As a result, setting the RHEL installation to **minimal-environment** installs only the necessary **s390utils-core** package and not the auxiliary **s390utils-base** package. If you want to use the **s390utils-base** package with a minimal RHEL installation, you must manually install the package after completing the RHEL installation or explicitly install **s390utils-base** using a kickstart file.

Bugzilla:1932480

4.2. RHEL FOR EDGE

Ignition support in RHEL for Edge Simplified images

With this enhancement, you can add an Ignition file to the Simplified Installer images by customizing your

blueprint. Both GUI and CLI have support for the Ignition customization. RHEL for Edge uses the Ignition provisioning utility to inject the user configuration into the images at an early stage of the boot process. On the first boot, Ignition reads its configuration either from a remote URL or a file embedded in the Simplified Installer image and applies that configuration into the image.

Jira:RHELPLAN-139659

Simplified Installer images can now be composed without the FDO customization section in the blueprint

Previously, to build a RHEL for Edge Simplified Installer image, you had to add details to the FIDO device onboarding (FDO) customization section. Otherwise, the image build would fail. With this update, the FDO customization in blueprints is now optional, and you can build RHEL for Edge Simplified Installer image with no errors.

Jira:RHELPLAN-139655

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.firewalld.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 **dnf offline-upgrade** command for offline updates on RHEL

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



IMPORTANT

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

Bugzilla:2131288


Applying advisory security filters to `dnf 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 **`dnf offline-upgrade`** command with advisory security filters (**`--advisory`**, **`--security`**, **`--bugfix`**, and other filters).

[Bugzilla:2139326](#)

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

With this enhancement, a new **`unload_plugins`** function has been added to the DNF API to allow plug-ins unloading.



IMPORTANT

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

[Bugzilla:2121662](#)

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:2150804](#)

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 2.6-17.el9 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:2046653](#)

`systemd` rebased to version 252

The **`systemd`** package has been upgraded to version 252. Notable changes include:

- You can specify the default timeout when waiting for device units to activate by using the **`DefaultDeviceTimeoutSec=`** option in **`system.conf`** and **`user.conf`** files.
- At shutdown, **`systemd`** now logs about processes blocking unmounting of file systems.
- You can now use drop-ins for transient units too.
- You can use size suffixes, such as K, M, G, T and others in the **`ConditionMemory=`** option.

- You can list automount points by using the **systemctl list-automounts** command.
- You can use the **systemd-logind** utility to stop an idle session after a preconfigured timeout by using the **StopIdleSessionSec=** option.
- The **systemd-udev** utility now creates the **infiniband by-path** and **infiniband by-ibdev** links for Infiniband verbs devices.
- The **systemd-tmpfiles** utility now gracefully handles the absent source of **C** copy.
- The **systemd-repart** utility now generates **dm-verity** partitions, including signatures.

[Bugzilla:2217931](#)

Updated **systemd-udev** assigns consistent network device names to InfiniBand interfaces

Introduced in RHEL 9, the new version of the **systemd** package contains the updated **systemd-udev** device manager. The device manager changes the default names of InfiniBand interfaces to consistent names selected by **systemd-udev**.

You can define custom naming rules for naming InfiniBand interfaces by following the [Renaming IPoB devices](#) procedure.

For more details of the naming scheme, see the **systemd.net-naming-scheme(7)** man page.

[Bugzilla:2136937](#)

4.5. INFRASTRUCTURE SERVICES

chrony rebased to version 4.3

The **chrony** suite has been updated to version 4.3. Notable enhancements over version 4.2 include:

- Added long-term quantile-based filtering of Network Time Protocol (NTP) measurements. You can enable this feature by adding the **maxdelayquant** option to the **pool**, **server**, or **peer** directive.
- Added the selection log to provide more information about **chronyd** selection of sources. You can enable the selection log by adding the **selection** option to the **log** directive.
- Improved synchronization stability when using the hardware timestamping and Pulse-Per-Second Hardware Clock (PHC) reference clocks.
- Added support for the system clock stabilization using a free-running stable clock, for example, Temperature Compensated Crystal Oscillator (TCXO), Oven-Controlled Crystal Oscillator (OCXO), or an atomic clock.
- Increased the maximum polling rate to 128 messages per second.

[Bugzilla:2133754](#)

frr rebased to version 8.3.1

The **frr** package for managing dynamic routing stack has been updated to version 8.3.1. Notable changes over version 8.2.2 include:

- Added a new set of commands to interact with the Border Gateway Protocol (BGP):

- the **set as-path replace** command to replace the Autonomous System (AS) path attribute of a BGP route with a new value.
- the **match peer** command to match a specific BGP peer or group when configuring a BGP route map.
- the **ead-es-frag evi-limit** command to set a limit on the number of Ethernet A-D per EVI fragments that can be sent in a given period of time in EVPN.
- the **match evpn route-type** command to take specific actions on certain types of EVPN routes, such as route-target, route-distinguisher, or MAC/IP routes.
- Added the **show thread timers** command in the VTYSH command-line interface for interacting with FRR daemons.
- Added the **show ip ospf reachable-routers** command to display a list of routers that are currently reachable through the OSPF protocol.
- Added new commands to interact with the Protocol Independent Multicast (PIM) daemon:
 - the **debug igmp trace detail** command to enable debugging for Internet Group Management Protocol (IGMP) messages with detailed tracing.
 - the **ip pim passive** command to configure the interface as passive, not sending PIM messages.
- Added new outputs for the **show zebra** command, such as ECMP, EVPN, MPLS statuses.
- Added the **show ip nht mrrib** command to the ZEBRA component to display multicast-related information from the **mrrib** table in the kernel.

[Bugzilla:2129731](#)

vsftpd rebased to version 3.0.5

The Very Secure FTP Daemon (**vsftpd**) provides a secure method of transferring files between hosts. The **vsftpd** package has been updated to version 3.0.5. Notable changes and enhancements include the following SSL modernizations:

- By default, the **vsftpd** utility now requires the use of TLS version 1.2 or later for secure connections.
- The **vsftpd** utility is now compatible with the latest FileZilla client.

[Bugzilla:2018284](#)

The frr package now contains targeted SELinux policy

Due to the fast development of the **frr** package for managing dynamic routing stack, new features and access vector cache (AVC) issues arose frequently. With this enhancement, the SELinux rules are now packaged together with FRR to address any issues faster. SELinux adds an additional level of protection to the package by enforcing mandatory access control policies.

[Bugzilla:2129743](#)

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:2044132

The **systemd-sysusers** utility is available in the **chrony**, **dhcpcd**, **radvd**, and **squid** packages

The **systemd-sysusers** utility creates system users and groups during package installation and removes them during a removal of the package. With this enhancement, the following packages contain the **systemd-sysusers** utility in their scriptlets:

- **chrony**,
- **dhcpcd**,
- **radvd**,
- **squid**.

Jira:RHELPLAN-136485

New **sync4l** 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 **sync4l** 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:2143264

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:2133815](#), [Bugzilla:2113925](#), [Bugzilla:2118786](#), [Bugzilla:2095829](#)

4.6. SECURITY

Libreswan rebased to 4.9

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

- Support for the **{left,right}pubkey=** options to the **addconn** and **whack** utilities
- KDF self-tests
- Show host's authentication key (**showhostkey**):
 - Support for ECDSA public keys
 - New **--pem** option to print 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
- The **pluto** IKE daemon:
 - Support for **maxbytes** and **maxpacket** counters

[Bugzilla:2128669](#)

OpenSSL rebased to 3.0.7

The OpenSSL packages have been rebased to version 3.0.7, which contains various bug fixes and enhancements. Most notably, the default provider now includes the **RIPEMD160** hash function.

[Bugzilla:2129063](#)

libssh now supports smart cards

You can now use smart cards through Public-Key Cryptography Standard (PKCS) #11 Uniform Resource Identifier (URI). As a result, you can use smart cards with the **libssh** SSH library and with applications that use **libssh**.

[Bugzilla:2026449](#)

libssh rebased to 0.10.4

The **libssh** library, which implements the SSH protocol for secure remote access and file transfer between machines, has been updated to version 0.10.4.

New features:

- Support for OpenSSL 3.0 has been added.
- Support for smart cards has been added.

- Two new configuration options **IdentityAgent** and **ModuliFile** have been added.

Other notable changes include:

- OpenSSL versions older than 1.0.1 are no longer supported
- By default, Digital Signature Algorithm (DSA) support has been disabled at build time.
- The SCP API has been deprecated.
- The **pubkey** and **privatekey** APIs have been deprecated.

[Bugzilla:2068475](#)

SELinux user-space packages updated to 3.5

The SELinux user-space packages **libselinux**, **libsepol**, **libsemanage**, **checkpolicy**, **mcstrans**, and **policycoreutils**, which includes the **sepolicy** utility, have been updated to version 3.5. Notable enhancements and bug fixes include:

- The **sepolicy** utility:
 - Added missing booleans to man pages
 - Several Python and GTK updates
- Added a workaround to **libselinux** that reduces heap memory usage by the **PCRE2** library
- The **libsepol** package:
 - Rejects attributes in type AV rules for kernel policies
 - No longer writes empty class definitions, which allows simpler round-trip tests
 - Stricter policy validation
- The **fixfiles** script unmounts temporary bind mounts on the **SIGINT** signal
- Many code and spelling bugs fixed
- Removed dependency on the deprecated Python module **distutils** and the installation using PIP
- The **semodule** option **--rebuild-if-modules-changed** renamed to **--refresh**
- Translation updated for generated descriptions and improved handling of unsupported languages
- Fixed many static code analysis bugs, fuzzer problems, and compiler warnings

[Bugzilla:2145224](#), [Bugzilla:2145228](#), [Bugzilla:2145229](#), [Bugzilla:2145226](#), [Bugzilla:2145230](#), [Bugzilla:2145231](#)

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:2159286](#)

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:

- New CIS RHEL9 profiles
- Deprecation of rule **account_passwords_pam_faillock_audit** in favor of **accounts_passwords_pam_faillock_audit**

[Bugzilla:2158405](#)

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:2122325](#)

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

Rules in **scap-security-guide** for checking and remediating ownership, group ownership, and permissions of Rsyslog log files are now also compatible with 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:2169414](#)

Keylime rebased to 6.5.2

The **keylime** packages have been rebased to upstream version - keylime-6.5.2-5.el9. This version contains various enhancements and bug fixes, most notably the following:

- Addressed vulnerability [CVE-2022-3500](#)
- The Keylime agent no longer fails IMA attestation when one scripts is executed quickly after another [RHBZ#2138167](#)
- Fixed segmentation fault in the **/usr/share/keylime/create_mb_refstate** script [RHBZ#2140670](#)
- Registrar no longer crashes during EK validation when the **require_ek_cert** option is enabled [RHBZ#2142009](#)

[Bugzilla:2150830](#)

Clevis accepts external tokens

With the new **-e** option introduced to the Clevis automated encryption tool, you can provide an external token ID to avoid entering your password during **cryptsetup**. This feature makes the configuration process more automated and convenient, and is useful particularly for packages such as **stratis** that use Clevis.

[Bugzilla:2126533](#)

Rsyslog TLS-encrypted logging now supports multiple CA files

With the new **NetstreamDriverCaExtraFiles** directive, you can specify a list of additional certificate authority (CA) files for TLS-encrypted remote logging. Note that the new directive is available only for the **ossli** (OpenSSL) Rsyslog network stream driver.

[Bugzilla:2124849](#)

Rsyslog privileges are limited

The privileges of the Rsyslog log processing system are now limited to only the privileges explicitly required by Rsyslog. This minimizes security exposure in case of a potential error in input resources, for example, a networking plugin. As a result, Rsyslog has the same functionality but does not have unnecessary privileges.

[Bugzilla:2127404](#)

SELinux policy allows Rsyslog to drop privileges at start

Because the privileges of the Rsyslog log processing system are now more limited to minimize security exposure ([RHBZ#2127404](#)), the SELinux policy has been updated to allow the **rsyslog** service to drop privileges at start.

[Bugzilla:2151841](#)

Tang now uses systemd-sysusers

The Tang network presence server now adds system users and groups through the **systemd-sysusers** service instead of shell scripts containing **useradd** commands. This simplifies checking of the system user list, and you can also override definitions of system users by providing **sysuser.d** files with higher priority.

[Bugzilla:2095474](#)

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:2110314

SELinux now confines mptcpd and udftools

With this update of the **selinux-policy** packages, SELinux confines the following services:

- **mptcpd**
- **udftools**

Bugzilla:1972222

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.

Jira:RHEL-192

GnuTLS can add and remove padding during decryption and encryption

The implementation of certain protocols requires PKCS#7 padding during decryption and encryption. The **gnutls_cipher_encrypt3** and **gnutls_cipher_decrypt3** block cipher functions have been added to GnuTLS to transparently handle padding. As a result, you can now use these functions in combination with the **GNUTLS_CIPHER_PADDING_PKCS7** flag to automatically add or remove padding if the length of the original plaintext is not a multiple of the block size.

[Bugzilla:2084161](#)

NSS no longer support RSA keys shorter than 1023 bits

The update of the Network Security Services (NSS) libraries changes the minimum key size for all RSA operations from 128 to 1023 bits. This means that NSS no longer perform the following functions:

- Generate RSA keys shorter than 1023 bits.
- Sign or verify RSA signatures with RSA keys shorter than 1023 bits.
- Encrypt or decrypt values with RSA key shorter than 1023 bits.

[Bugzilla:2091905](#)

The Extended Master Secret TLS Extension is now enforced on FIPS-enabled systems

With the release of the [RHSA-2023:3722](#) advisory, the TLS **Extended Master Secret** (EMS) extension (RFC 7627) is mandatory for TLS 1.2 connections on FIPS-enabled RHEL 9 systems. This is in accordance with FIPS-140-3 requirements. TLS 1.3 is not affected.

Legacy clients that do not support EMS or TLS 1.3 now cannot connect to FIPS servers running on RHEL 9. Similarly, RHEL 9 clients in FIPS mode cannot connect to servers that only support TLS 1.2 without EMS. This in practice means that these clients cannot connect to servers on RHEL 6, RHEL 7 and non-RHEL legacy operating systems. This is because the legacy 1.0.x versions of OpenSSL do not support EMS or TLS 1.3.

In addition, connecting from a FIPS-enabled RHEL client to a hypervisor such as VMWare ESX now fails with a **Provider routines::ems not enabled** error if the hypervisor uses TLS 1.2 without EMS. To work around this problem, update the hypervisor to support TLS 1.3 or TLS 1.2 with the EMS extension. For

VMWare vSphere, this means version 8.0 or later.

For more information, see [TLS Extension "Extended Master Secret" enforced with Red Hat Enterprise Linux 9.2](#).

[Bugzilla:2188046](#), [Bugzilla:2218721](#)

4.7. NETWORKING

NetworkManager rebased to version 1.42.2

The **NetworkManager** packages have been upgraded to upstream version 1.42.2, which provides a number of enhancements and bug fixes over the previous version:

- Ethernet bonds support source load balancing.
- NetworkManager can manage connections on the **loopback** device.
- Support for IPv4 equal-cost multi-path (ECMP) routes was added.
- Support for **802.1ad** tagging in Virtual Local Area Networks (VLANs) connections was added.
- The **nmtui** application supports Wi-Fi WPA-Enterprise, Ethernet with 802.1X authentication, and MACsec connection profiles.
- NetworkManager rejects DHCPv6 leases if all addresses fail IPv6 duplicate address detection (DAD).

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

[Bugzilla:2134897](#)

Introduction of the **weight** property in ECMP routing with NetworkManager

With this update, RHEL 9 supports a new property **weight** when defining IPv4 Equal-Cost Multi-Path (ECMP) routes. You can configure multipath routing using NetworkManager to load-balance and stabilize network traffic. This allows for multiple paths to be used for data transmission between two nodes, which improves the network efficiency and provides redundancy in the event of a link failure. Conditions for using the **weight** property include:

- The valid values are 1-256.
- Define multiple next-hop routes as single-hop routes with the **weight** property.
- If you do not set **weight**, NetworkManager cannot merge the routes into an ECMP route.

[Bugzilla:2081302](#)

NetworkManager update brings improved flexibility for DNS configuration across multiple networks

With this update, you can use the existing **[global-dns]** section in the **/etc/NetworkManager/NetworkManager.conf** file to configure DNS options without specifying the **nameserver** value in the **[global-dns-domain-*)** section. This enables you to configure DNS options in the **/etc/resolv.conf** file while still relying on the DNS servers provided by the network connection for

actual DNS resolution. As a result, the feature makes it easier and more flexible to manage your DNS settings when connecting to different networks with different DNS servers. Especially when you use the **/etc/resolv.conf** file to configure DNS options.

[Bugzilla:2019306](#)

NetworkManager now supports a new **vlan.protocol** property

With this update, the **vlan** interface type now accepts a new **protocol** property. The property type is string. The accepted values are either **802.1Q** (default), or **802.1ad**. The new property specifies which VLAN protocol controls the tag identifier for encapsulation.

[Bugzilla:2128809](#)

NetworkManager now allows VLAN configuration over unmanaged interface

With this enhancement, you can use an unmanaged networking interface as a base interface when configuring virtual LAN (VLAN) with NetworkManager. As a result, the VLAN base interface remains intact unless changed explicitly through the **nmcli device set enp1s0 managed true** command or other API of NetworkManager.

[Bugzilla:2110307](#)

Configuring Multipath TCP using NetworkManager is now fully supported

With this update, the NetworkManager utility provides you with the Multipath TCP (MPTCP) functionality. You can use **nmcli** commands to control MPTCP and make its settings persistent.

For more information, see:

- [Understanding Multipath TCP: High availability for endpoints and the networking highway of the future](#)
- [RFC 8684: TCP Extensions for Multipath Operation with Multiple Addresses](#)
- [Permanently configuring multiple paths for MPTCP applications](#)

[Bugzilla:2029636](#)

The NetworkManager utility now supports activating connections on the **loopback** interface

Administrators can manage the **loopback** interface to:

- Add extra IP addresses to the **loopback** interface
- Define DNS configuration
- Define a special route, which does not bind to an interface
- Define a route rule, which is not interface-related
- Change Maximum Transmission Unit (MTU) size of the **loopback** interface

[Bugzilla:2073512](#)

The **balance-slb** bonding mode is now supported

The new **balance-slb** bonding mode Source load balancing requires no switch configuration. The

balance-slb divides traffic on the source ethernet address using **xmit_hash_policy=vlan+srcmac**, and NetworkManager adds necessary **nftables** rules for traffic filtering. As a result, you can now create bond profiles with the **balance-slb** option enabled by using NetworkManager.

[Bugzilla:2128216](#)

firewalld rebased to version 1.2

The **firewalld** package has been upgraded to version 1.2, which provides multiple enhancements. Notable changes include:

- Support for new services (for example Kodi JSON-RPC, EventServer, netdata, IPFS)
- Fail-safe mode to ensure that the system remains protected and that network communication is not disrupted if the **firewalld** service encounters an error during its startup
- Tab-completion in command-line (CLI) for some of the **firewalld** policy commands

[Bugzilla:2125371](#)

The firewalld now supports the startup failsafe mechanism

With this enhancement, **firewalld** will fall back to failsafe defaults in case of a startup failure. This feature protects the host in case of invalid configurations or other startup issues. As a result, even if the user configuration is invalid, hosts running **firewalld** are now startup failsafe.

[Bugzilla:2077512](#)

conntrack-tools rebased to version 1.4.7

The **conntrack-tools** package has been upgraded to version 1.4.7, which provides multiple bug fixes and enhancements. Notable changes include:

- Adds the **IPS_HW_OFFLOAD** flag, which specifies offloading of a **conntrack** entry to the hardware
- Adds **clash_resolve** and **chaintoolong** statistical counters
- Supports filtering events by IP address family
- Accepts yes or no as synonyms to on or off in the **conntrackd.conf** file
- Supports user space helper auto-loading upon daemon startup. Users do not have to manually run the **nfct add helper** commands
- Removes the **-o userspace** command option and always tags user space triggered events
- Logs external inject problems as warning only
- Ignores conntrack ID when looking up cache entries to allow for stuck old ones to be replaced
- Fixes broken parsing of IPv6 **M-SEARCH** requests in the **ssdp cthelper** module
- Eliminates the need for lazy binding technique in the **nfct** library
- Sanitizes protocol value parsing, catch invalid values

[Bugzilla:2132398](#)

The **nmstate** API now supports IPv6 link-local addresses as DNS servers

With this enhancement, you can use the **nmstate** API to set IPv6 link-local addresses as DNS servers. Use the **<link-local_address>%<interface>** format, for example:

```
dns-resolver:
  config:
    server:
      - fe80::deef:1%enp1s0
```

[Bugzilla:2095207](#)

The **nmstate** API now supports MPTCP flags

This update enhances the **nmstate** API with support for MultiPath TCP (MPTCP) flags. As a result, you can use **nmstate** to set MPTCP address flags on interfaces with static or dynamic IP addresses.

[Bugzilla:2120473](#)

The **min-mtu** and **max-mtu** properties added to MTU on all interfaces

Previously, an exception message was not clear enough to understand the supported MTU ranges. This update introduces the **min-mtu** and **max-mtu** properties to all interfaces. As a result, **nmstate** will indicate the supported MTU range when the desired MTU is out of range.

[Bugzilla:2044150](#)

NetworkManager now allows VLAN configuration over unmanaged interface

With this enhancement, you can use an unmanaged networking interface as a base interface when configuring virtual LAN (VLAN) with NetworkManager. As a result, the VLAN base interface remains intact unless changed explicitly through the **nmcli device set *enp1s0* managed true** command or other API of NetworkManager.

[Bugzilla:2058292](#)

The **balance-slb** bonding mode is now supported

The new **balance-slb** bonding mode Source load balancing requires no switch configuration. The **balance-slb** divides traffic on the source Ethernet address using **xmit_hash_policy=vlan+srcmac**, and NetworkManager adds necessary **nftables** rules for traffic filtering. As a result, you can now create bond profiles with the **balance-slb** option enabled by using NetworkManager.

[Bugzilla:2130240](#)

A new **weight** property in Nmstate

This update introduces the **weight** property in the Nmstate API and tooling suite. You can use **weight** to specify the relative weight of each path in the Equal Cost Multi-Path routes (ECMP) group. The weight is a number between 1 and 256. As a result, **weight** property in Nmstate provides greater flexibility and control over traffic distribution in an ECMP group.

[Bugzilla:2162401](#)

xdp-tools rebased to version 1.3.1

The **xdp-tools** packages have been upgraded to upstream version 1.3.1, which provides a number of enhancements and bug fixes over the previous version:

- The following utilities have been added:
 - **xdp-bench**: Performs XDP benchmarks on the receive side.
 - **xdp-monitor**: Monitors XDP errors and statistics using kernel trace points.
 - **xdp-trafficgen**: Generates and sends traffic through the XDP driver hook.
- The following features have been added to the **libxdp** library:
 - The **xdp_multiprog_xdp_frags_support()**, **xdp_program__set_xdp_frags_support()**, and **xdp_program__xdp_frags_support()** functions have been added to support loading programs with XDP **frags** support, a feature that is also known as **multibuffer XDP**.
 - The library performs proper reference counting when attaching programs to **AF_XDP** sockets. As a result, the application no longer has to manually detach XDP programs when using sockets. The **libxdp** library detaches the program now automatically when the program is no longer used.
 - The following functions have been added to the library:
 - **xdp_program__create()** for creating **xdp_program** objects
 - **xdp_program__clone()** for cloning an **xdp_program** reference
 - **xdp_program__test_run()** for running XDP programs through the **BPF_PROG_TEST_RUN** kernel API
 - When the **LIBXDP_BPF_FS_AUTOMOUNT** environment variable is set, the **libxdp** library now supports automatically mounting of a **bpffs** virtual file system if none is found. A subset of the library features can now also function when no **bpffs** is mounted.

Note that this version also changes the version number of the XDP dispatcher program that is being loaded on the network devices. This means that you can not use a previous and a new version of **libxdp** and **xdp-tools** at the same time. The **libxdp** 1.3 library will display old versions of the dispatcher, but not automatically upgrade them. Additionally, after loading a program with **libxdp** 1.3, older versions will not interoperate with the newer one.

[Bugzilla:2160066](#)

iproute rebased to version 6.1.0

The **iproute** package has been upgraded to version 6.1.0, which provides multiple bug fixes and enhancements. Notable changes include:

- Supports reading the **vdpa** device statistics
 - Illustration of statistics reading for the **virtqueue** data structure at index 1:


```
# vdpas dev vstats show vdpas-a qidx 1
vdpas-a:
vdpas-a: queue_type tx received_desc 321812 completed_desc 321812
```
 - Illustration of statistics reading for the **virtqueue data** structure at index 16:


```
# vdpas dev vstats show vdpas-a qidx 16
vdpas-a: queue_type control_vq received_desc 17 completed_desc 17
```

- Updates the corresponding manual pages

[Bugzilla:2155604](#)

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:2143850](#)

4.8. KERNEL

Kernel version in RHEL 9.2

Red Hat Enterprise Linux 9.2 is distributed with the kernel version 5.14.0-284.11.1.

[Bugzilla:2177782](#)

The 64k page size kernel is now available

In addition to the RHEL 9 for ARM kernel which supports 4k pages, Red Hat now offers an optional kernel package that supports 64k pages: **kernel-64k**.

The 64k page size kernel is a useful option for large datasets on ARM platforms. It enables better performance for some types of memory- and CPU-intensive operations.

You must choose page size on 64-bit ARM architecture systems at the time of installation. You can install **kernel-64k** only by Kickstart by adding the **kernel-64k** package to the package list in the **Kickstart** file.

For more information on installing **kernel-64k**, see [Performing an advanced RHEL 9 installation](#) .

[Bugzilla:2153073](#)

virtiofs support for kexec-tools enabled

This enhancement adds the **virtiofs** feature for **kexec-tools** by introducing the new option, **virtiofs myfs**, where **myfs** is a variable tag name to set in the **qemu** command line, for example, **-device vhost-user-fs-pci,tag=myfs**

The **virtiofs** file system implements a driver that allows a guest to mount a directory that has been exported on the host. By using this enhancement, you can save the virtual machine's **vmcore** dump file to:

- A **virtiofs** shared directory.
- The sub-directory, such as **/var/crash**, when the root file system is a **virtiofs** shared directory.
- A different **virtiofs** shared directory, when the virtual machine's root file system is a **virtiofs** shared directory.

[Bugzilla:2085347](#)

The **kexec-tools** package now adds improvements on remote **kdump** targets

With this enhancement, the **kexec-tools** package adds significant bug fixes and enhancements. The most notable changes include:

- Optimized memory consumption for **kdump** by enabling only the required network interfaces.
- Improved network efficiency for **kdump** in events of connection timeout failures.
The default wait time for a network to establish is 10 minutes maximum. This removes the need to pass **dracut** parameters, such as **rd.net.timeout.carrier** or **rd.net.timeout.dhcp** as a workaround to identify a carrier.

[Bugzilla:2076416](#)

BPF rebased to version 6.0

The Berkeley Packet Filter (BPF) facility has been rebased to Linux kernel version 6.0 with multiple enhancements. This update enables all the BPF features that depend on the BPF Type Format (BTF) for kernel modules. Such features include the usage of BPF trampolines for tracing, the availability of the Compile Once - Run Everywhere (CO-RE) mechanism, and several networking-related features. Furthermore, the kernel modules now contain debugging information, which means that you no longer need to install **debuginfo** packages to inspect the running modules.

For more information on the complete list of BPF features available in the running kernel, use the **bpftool feature** command.

[Jira:RHELPLAN-133650](#)

The **rtla** meta-tool adds the **osnoise** and **timerlat** tracers for improved tracing 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 precise 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 information about operating system noise.
- 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:2075216](#)

The **argparse** module of Tuna now supports configuring CPU sockets

With this enhancement, you can specify a specific CPU socket when you have multiple CPU sockets. You can view the help usage by using the **-h** on a subcommand, for example, **tuna show_threads -h**.

To configure a specific CPU socket, specify the **-S** option with each **tuna** command where you need to use CPU sockets:

```
tuna <command> [-S CPU_SOCKET_LIST]
```

For example, use **tuna show_threads -S 2,3** to view the threads or **tuna show_irqs -S 2,3** to view attached interrupt requests (IRQs).

As a result, this enhancement facilitates CPU usage based on CPU sockets without the need to specify each CPU individually.

[Bugzilla:2122781](#)

The output format for **cgroups** and **irqs** in **Tuna** is 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, the **cgroups** output now has an improved readable format that is adaptable to your terminal size.

[Bugzilla:2121517](#)

A new command line interface has been added to the **tuna** tool in real-time

This enhancement adds a new command line interface to the **tuna** tool, which is based on the **argparse** parsing module. With this update, you can now perform the following tasks:

- Change the attributes of the application and kernel threads.
- Operate on interrupt requests (IRQs) by name or number.
- Operate on tasks or threads by using the process identifier.
- Specify CPUs and sets of CPUs with the CPU or the socket number.

By using the **tuna -h** command, you can print the command line arguments and their corresponding options. For each command, there are optional arguments, which you can view with the **tuna <command> -h** command.

As a result, **tuna** now provides an interface with a more standardized menu of commands and options that is easier to use and maintain than the command line interface.

[Bugzilla:2062865](#)

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:2081325](#)

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

With this enhancement, you can specify a latency range for a single bucket at nanoseconds 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:2041637](#)

The NVMe/FC transport protocol enabled as the **kdump** storage target

The **kdump** mechanism now provides the support for Nonvolatile Memory Express (NVMe) over Fibre Channel (NVMe/FC) protocol as the dump target. With this update, you can configure **kdump** to save kernel crash dump files on NVMe/FC storage targets.

As a result, **kdump** can capture and save the **vmcore** file on **NVMe/FC** in the event of a kernel crash without **timeout** or **reconnect** errors.

For more information on NVMe/FC configuration, see [Managing storage devices](#)

[Bugzilla:2080110](#)

The **crash-utility** tool has been rebased to version 8.0.2

The **crash-utility**, which analyzes an active system state or after a kernel crash, has been rebased to version 8.0.2. The notable change includes adding support for **multiqueue(blk-mq)** devices. By using the **dev -d** or **dev -D** command, you can display the disk I/O statistics for **multiqueue(blk-mq)** devices.

[Bugzilla:2119685](#)

openssl-ibmca rebased to version 2.3.1

The dynamic OpenSSL engine and provider for IBMCA on 64-bit IBM Z architecture have been rebased to upstream version 2.3.1. Users of RHEL 9 are recommended to use the OpenSSL *provider* to ensure compatibility with future updates of OpenSSL. The *engine* functionality has been deprecated in OpenSSL version 3.

[Bugzilla:2110378](#)

Secure Execution guest dump encryption with customer keys

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

[Bugzilla:2044204](#)

The TSN protocol for real-time has been enabled on the ADL-S platform

With this enhancement, the IEEE Time Sensitive Networking (TSN) specification enables time synchronization and deterministic processing of real-time workloads over the network on Intel Alder Lake S (ADL-S) platform. It supports the following network devices:

- A discrete 2.5GbE MAC-PHY combo with TSN support: Intel® i225/i226
- An integrated 2.5GbE MAC in the SOC with 3rd party PHY chips from Marvell, Maxlinear and TI covering the 1GbE and 2.5GbE speed, is available on select **skus** and SOCs.

With the TSN protocol, you can manage deterministic applications scheduling, preemption, and accurate time synchronization type workloads in embedded implementations. These implementations need dedicated, specialized, and proprietary networks, while workloads run on standard Ethernet, Wi-Fi, and 5G networks.

As a result, TSN provides improved capabilities for:

- Hardware: Intel based systems used for implementing real-time workloads in IoT
- Deterministic and time sensitive applications

Bugzilla:2100606

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:2104468

Option to write data for **gnss** module is now available

This update provides the option of writing data to the **gnss** receiver. Previously, **gnss** was not fully configurable. With this enhancement, all **gnss** functions are now available.

Bugzilla:2111048

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 Secure Boot CA 3** certificate that you need to load into the HMC is available at [Product Signing Keys](#).

Bugzilla:2190123

4.9. FILE SYSTEMS AND STORAGE

nvme-cli rebased to version 2.2.1

The **nvme-cli** packages have been upgraded to version 2.2.1, which provide multiple bug fixes and enhancements. Notable changes include:

- Added the new **nvme show-topology** command, which displays the topology of all NVMe subsystems.
- Dropped the **libuuid** dependency.
- The **uint128** data fields are displayed correctly.
- Updated the **libnvme** dependency to version 1.2.

Bugzilla:2139753

libnvme rebased to version 1.2

The **libnvme** packages have been upgraded to version 1.2, which provide multiple bug fixes and enhancements. The most notable change is a dropped dependency of the **libuuid** library.

[Bugzilla:2139752](#)

Stratis enforces consistent block size in pools

Stratis now enforces a consistent block size in pools to address potential edge case problems that can occur when mixed block size devices exist within a pool. With this enhancement, users can no longer create a pool or add new devices that have a different block size from the existing devices in the pool. As a result, there is a reduced risk of pool failure.

[Bugzilla:2039957](#)

Support for existing disk growth within the Stratis pool

Previously, when a user added new disks to the RAID array, the size of the RAID array would generally increase. However, in all cases, Stratis ignored the increase in size and continued to use only the space that was available on the RAID array when it was first added to the pool. As a result, Stratis was unable to identify the new device, and users could not increase the size of the pool.

With this enhancement, Stratis now identifies any pool device members that have expanded in size. As a result, users can now issue a command to expand the pool based on their requirements.

Stratis now supports the growth of existing disks within its pool, in addition to the existing feature of growing the pool by adding new disks.

[Bugzilla:2039955](#)

Improved functionality of the **lvreduce** command

With this enhancement, when the logical volume (LV) is active, the **lvreduce** command checks if reducing the LV size would damage any file system present on it. If a file system on the LV requires reduction, and the **lvreduce resizefs** option has not been enabled, then the LV will not be reduced.

Additionally, new options are now available to control the handling of file systems while reducing an LV. These options provide users with greater flexibility and control when using the **lvreduce** command.

[Bugzilla:1878893](#)

Direct I/O alignment information for **statx** was added

This update introduces a new mask value, "**STATX_DIOALIGN**", to the **statx(2)** call. When this value is set in the **stx_mask** field, it requests **stx_dio_mem_align** and **stx_dio_offset_align** values, which indicate the required alignment (in bytes) for user memory buffers and file offsets and I/O segment lengths for direct I/O (O_DIRECT) on this file, respectively. If direct I/O is not supported on the file, both values will be 0. This interface is now implemented for block devices as well as for files on the xfs and ext4 filesystems in RHEL9.

[Bugzilla:2150284](#)

NFSv4.1 session trunking discovery

With this update, the client can use multiple connections to the same server and session, resulting in faster data transfer. When an NFS client mounts a multi-homed NFS server with different IP addresses, only one connection is used by default, ignoring the rest. To improve performance, this update adds support for the **trunkdiscovery** and **max_connect** mount options, which enable the client to test each connection and associate multiple connections with the same NFSv4.1+ server and session.

[Bugzilla:2066372](#)

NFS IO sizes can now be set as a multiples of **PAGE_SIZE** for TCP and RDMA

This update allows users to set NFS IO sizes as a multiples of **PAGE_SIZE** for TCP and RDMA connections. This offers greater flexibility in optimizing NFS performance for some architectures.

Bugzilla:2107347

nfsrahead has been added to RHEL 9

With the introduction of the **nfsrahead** tool, you can use it to modify the **readahead** value for NFS mounts, and thus affect the NFS read performance.

Bugzilla:2143747

4.10. 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:2116295

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:2112270, Bugzilla:2159454

4.11. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

Python 3.11 available in RHEL 9

RHEL 9.2 introduces Python 3.11, provided by the new package **python3.11** and a suite of packages built for it, as well as the **ubi9/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 on the same system.

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

```
# dnf install python3.11
# dnf 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 Python 3.11 will have a shorter life cycle than Python 3.9, which is the default Python implementation in RHEL 9; see [Red Hat Enterprise Linux Application Streams Life Cycle](#) .

[Bugzilla:2127923](#)

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

The updated **Node.js 18.14** 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:2178088](#)

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=reword:<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:2139379](#)

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:2139383](#)

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:

```
# dnf module install nginx:1.22
```

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:2096174

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 2.9.3.

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:2143211](#)

New packages: tomcat

RHEL 9.2 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:2160511](#)

A new module stream: postgresql:15

RHEL 9.2 introduces **PostgreSQL 15** as the **postgresql:15** module stream. **PostgreSQL 15** 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:

```
# dnf module install postgresql:15
```

If you want to upgrade from an earlier **postgresql** stream within RHEL 9, migrate your **PostgreSQL** data as described in [Migrating to a RHEL 9 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:2128410](#)

4.12. COMPILERS AND DEVELOPMENT TOOLS

openblas rebased to version 0.3.21

The OpenBLAS library has been updated to version 0.3.21. This update includes performance optimization patches for the IBM POWER10 platform.

[Bugzilla:2112099](#)

A new module stream: swig:4.1

RHEL 9.2 introduces the Simplified Wrapper and Interface Generator (SWIG) version 4.1 as the **swig:4.1** module stream available in the CodeReady Linux Builder (CRB) repository. Note that packages included in the CodeReady Linux Builder repository are unsupported.

Compared to **SWIG 4.0** released in RHEL 9.0, **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:

1. Enable the [CodeReady Linux Builder \(CRB\) repository](#).
2. Install the module stream:

```
# dnf module install swig:4.1
```

[Bugzilla:2139101](#)

New package: **jmc** in the CRB repository

RHEL 9.2 introduces the JDK Mission Control (JMC) profiler for HotSpot JVMs version 8.2.0, available as the **jmc** package in the CodeReady Linux Builder (CRB) repository for the AMD and Intel 64-bit architectures.

To install JMC, you must first enable the [CodeReady Linux Builder \(CRB\) repository](#).

Note that packages included in the CRB repository are unsupported.

[Bugzilla:2122401](#)

OpenJDK service attributes now available in FIPS mode

Previously, cryptographic services and algorithms available for OpenJDK in FIPS mode were too strictly filtered and resulted in unavailable service attributes. With this enhancement, these service attributes are now available in FIPS mode.

[Bugzilla:2186803](#)

Performance Co-Pilot rebased to version 6.0

Performance Co-Pilot (PCP) has been updated to version 6.0. Notable improvements include:

1. Version 3 PCP archive support:

This includes support for instance domain change-deltas, Y2038-safe timestamps, nanosecond-precision timestamps, arbitrary timezones support, and 64-bit file offsets used throughout for larger (beyond 2GB) individual volumes.

This feature is currently opt-in via the **PCP_ARCHIVE_VERSION** setting in the **/etc/pcp.conf** file.

Version 2 archives remain the default.

2. Only OpenSSL is used throughout PCP. Mozilla NSS/NSPR use has been dropped: This impacts **libpcp**, **PMAPI** clients and **PMCD** use of encryption. These elements are now configured and used consistently with **pmproxy** HTTPS support and **redis-server**, which were both already using OpenSSL.
3. New nanosecond precision timestamp **PMAPI** calls for **PCP** library interfaces that make use of timestamps. These are all optional, and full backward compatibility is preserved for existing tools.
4. The following tools and services have been updated:

pcp2elasticsearch

Implemented authentication support.

pcp-dstat

Implemented support for the **top-alike** plugins.

pcp-htop

Updated to the latest stable upstream release.

pmseries

Added **sum**, **avg**, **stdev**, **nth_percentile**, **max_inst**, **max_sample**, **min_inst** and **min_sample** functions.

pmdabpf

Added CO-RE (Compile Once - Run Everywhere) modules and support for AMD64, Intel 64-bit, 64-bit ARM, and IBM Power Systems.

pmdabpftrace

Moved example autostart scripts to the **/usr/share** directory.

pmdadenki

Added support for multiple active batteries.

pmdalinux

Updates for the latest **/proc/net/netstat** changes.

pmdaopenvswitch

Added additional interface and coverage statistics.

pmproxy

Request parameters can now be sent in the request body.

pmieconf

Added several **pmie** rules for Open vSwitch metrics.

pmlogger_farm

Added a default configuration file for farm loggers.

pmlogger_daily_report

Some major efficiency improvements.

[Bugzilla:2117074](#)

grafana rebased to version 9.0.9

The **grafana** package has been rebased to version 9.0.9. Notable changes include:

- The time series panel is now the default visualization option, replacing the graph panel
- New heatmap panel
- New Prometheus and Loki query builder
- Updated Grafana Alerting
- Multiple UI/UX and performance improvements
- License changed from Apache 2.0 to GNU Affero General Public License (AGPL)

The following are offered as opt-in experimental features:

- New bar chart panel
- New state timeline panel
- New status history panel
- New histogram panel

For more information, see: [What's new in Grafana v9.0](#) and [What's new in Grafana v8.0](#).

[Bugzilla:2116847](#)

grafana-pcp rebased to version 5.1.1

The **grafana-pcp** package has been rebased to version 5.1.1. Notable changes include:

Query editor

added buttons to disable rate conversion and time utilization conversion.

Redis

removed the deprecated **label_values(metric, label)** function.

Redis

fixed the network error for metrics with many series (requires Performance Co-Pilot v6+).

Redis

set the **pmproxy** API timeout to 1 minute.

[Bugzilla:2116848](#)

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 9.2 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:

```
# dnf 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:2110583](#)

The updated GCC compiler is now available for RHEL 9.2

The system GCC compiler, version 11.3.1, has been updated to include numerous bug fixes and enhancements available in the upstream GCC.

The GNU Compiler Collection (GCC) provides tools for developing applications with the C, C++, and Fortran programming languages.

For usage information, see [Developing C and C++ applications in RHEL 9](#) .

[Bugzilla:2117632](#)

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:2118567](#)

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:2123900](#)

Go Toolset rebased to version 1.19.6

Go Toolset has been updated to version 1.19.6. 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:2175173](#)

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:2157982

4.13. 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

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:1507035

IdM now supports the **min_lifetime** parameter

With this enhancement, the **min_lifetime** parameter has been added to the **/etc/gssproxy/*.conf** file. The **min_lifetime** parameter triggers the renewal of a service ticket in case its remaining lifetime is lower than this value.

By default its value is 15 seconds. For network volume clients such as NFS, to reduce the risk of losing access in case the KDC is momentarily unavailable, set this value to 60 seconds.

Bugzilla:2184333

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

IdM now enforces the presence of the MS-PAC structure in Kerberos tickets

Starting with RHEL 9.2, to increase security, Identity Management (IdM) and MIT Kerberos now enforce the presence of the Privilege Attribute Certificate (MS-PAC) structure in the Kerberos tickets issued by the RHEL IdM Kerberos Distribution Center (KDC).

In November 2022, in response to CVE-2022-37967, Microsoft introduced an extended signature that is calculated over the whole MS-PAC structure rather than over the server checksum. Starting with RHEL 9.2, the Kerberos tickets issued by IdM KDC now also contain the extended signature.



NOTE

The presence of the extended signature is not yet enforced in IdM.

Jira:RHELPLAN-159146

New realm configuration template for KDC enabling FIPS 140-3-compliant key encryption

This update provides a new, **EXAMPLE.COM**, example realm configuration in the **/var/kerberos/krb5kdc/kdc.conf** file. It brings two changes:

- The FIPS 140-3-compliant **AES HMAC SHA-2** family is added to the list of supported types for key encryption.
- The encryption type of the KDC master key is switched from **AES 256 HMAC SHA-1** to **AES 256 HMAC SHA-384**.



WARNING

This update is about standalone MIT realms. Do not change the Kerberos Distribution Center (KDC) configuration in RHEL Identity Management.

Using this configuration template is recommended for new realms. The template does not affect any realm already deployed. If you are planning to upgrade the configuration of your realm according to the template, consider the following points:

For upgrading the master key, changing the setting in the KDC configuration is not enough. Follow the process described in the MIT Kerberos documentation: <https://web.mit.edu/kerberos/krb5-1.20/doc/admin/database.html#updating-the-master-key>

Adding the **AES HMAC SHA-2** family to the supported types for key encryption is safe at any point

because it does not affect existing entries in the KDC. Keys will be generated only when creating new principals or when renewing credentials. Note that keys of this new type cannot be generated based on existing keys. To make these new encryption types available for a certain principal, its credentials have to be renewed, which means renewing keytabs for service principals too.

The only case where principals should not feature an **AES HMAC SHA-2** key is the Active Directory (AD) cross-realm ticket-granting ticket (TGT) ones. Because AD does not implement RFC8009, it does not use the **AES HMAC SHA-2** encryption types family. Therefore, a cross-realm TGS-REQ using an **AES HMAC SHA-2**-encrypted cross-realm TGT would fail. The best way to keep the MIT Kerberos client from using **AES HMAC SHA-2** against AD is to not provide **AES HMAC SHA-2** keys for the AD cross-realm principals. To do so, ensure that you create the cross-realm TGT entries with an explicit list of key encryption types that are all supported by AD:

```
kadmin.local <<EOF
add_principal +requires_preauth -e aes256-cts-hmac-sha1-96,aes128-cts-hmac-sha1-96 -pw
[password] krbtgt/[MIT realm]@[AD realm]
add_principal +requires_preauth -e aes256-cts-hmac-sha1-96,aes128-cts-hmac-sha1-96 -pw
[password] krbtgt/[AD realm]@[MIT realm]
EOF
```

To ensure the MIT Kerberos clients use the **AES HMAC SHA-2** encryption types, you must also set these encryption types as **permitted** in both the client and the KDC configuration. On RHEL, this setting is managed by the crypto-policy system. For example, on RHEL 9, hosts using the **DEFAULT** crypto-policy allow **AES HMAC SHA-2** and **AES HMAC SHA-1** encrypted tickets, while hosts using the **FIPS** crypto-policy only accept **AES HMAC SHA-2** ones.

[Bugzilla:2068535](#)

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:2126640](#), [Bugzilla:2142805](#)

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:<l>{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:2131993](#)

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:2143224](#)

Red Hat IdM and Certificate System now support the EST protocol

Enrollment over Secure Transport (EST) is a new Certificate System subsystem feature that is specified in RFC 7030 and it is used to provision certificates from a Certificate Authority (CA). EST implements the server side of the operation, such as **/getcacerts**, **/simpleenroll**, and **/simplereenroll**.

Note that Red Hat supports both EST and the original Simple Certificate Enrollment Protocol (SCEP) in Certificate System.

[Bugzilla:1849834](#)

Automatic purging of expired certificates

This update adds an automatic mechanism to purge expired certificates and request records from the database. You can enable this feature based on certain policies, such as search size limit, search time limit, and retention time.

To remove records safely, the CA needs to use the Random Certificate Serial Numbers v1 (RSNv3) to generate the certificate serial numbers and enrollment or renewal request IDs. The CA provides a pruning job that removes the following:

- Certificates that have expired for some time.
- Completed requests corresponding to the expired certificates.
- Incomplete requests that have been idle for some time.

You need to schedule this job to run regularly to remove a certain number of records each time it runs. The remaining records will be removed in the subsequent runs. For large deployments, you can distribute the job among the servers in the cluster.

[Bugzilla:1883477](#)

Enhance negative cache usage

This update improves the SSSD performance for lookups by Security Identifier (SID). It now stores non-existing SIDs in the negative cache for individual domains and requests the domain that the SID belongs to.

[Bugzilla:1766490](#)

ACME now supports automatically removing expired certificates

Previously, the Automated Certificate Management Environment (ACME) service in Identity Management (IdM) did not remove expired certificates from the certificate authority (CA). Because ACME issues short-lived, 90-day certificates, expired certificates accumulated in the CA, which affected performance.

With this enhancement, ACME can now automatically remove expired certificates at specified intervals.

Removing expired certificates is disabled by default. To enable it, enter:

```
# ipa-acme-manage pruning --enable --cron "0 0 1 * **"
```

This removes expired certificates on the first day of every month at midnight.



NOTE

Expired certificates are removed after their retention period. By default, this is 30 days after expiry.

For more details, see the **ipa-acme-manage(1)** man page.

[Bugzilla:2162677](#)

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](#)

Directory Server now supports extended logging of search operations

Previously, records in the access log did not show why some search operations had a very big **etime** value. With this release, you can enable logging of statistics such as a number of index lookups (database read operations) and overall duration of index lookups per each search operation. These statistical records can help to analyze why the **etime** value can be so resource expensive.

[Bugzilla:1859271](#)

The NUNC_STANS error logging level was replaced by the new **1048576** logging level

Previously, you could not easily debug password policy issues. With the new **1048576** logging level for the error log, you can now check the following password policy information:

- Which local policy rejects or allows a password update.
- The exact syntax violation.

[Bugzilla:2057070](#)

Directory Server introduces the security log

To properly track issues over time, Directory Server now has a specialized log that maintains security data. The security log does not rotate quickly and consumes less disk resources in comparison to the access log that has all the information, but requires expensive parsing to get the security data.

The new server log records security events such as authentication events, authorization issues, DoS/TCP attacks, and other events.

Directory Server stores the security log in the **/var/log/dirsrv/slapd-*instance_name*** directory along with other log files.

[Bugzilla:2093981](#)

Directory Server now can compress archived log files

Previously, archived log files were not compressed. With this release, you can enable access, error, audit, audit fail log, security log files compression to save disk space. Note that only security log file compression is enabled by default.

Use the following new configuration attributes in the **cn=config** entry to manage the compression:

- **nsslapd-accesslog-compress** for the access log
- **nsslapd-errorlog-compress** for the error log
- **nsslapd-auditlog-compress** for the audit log
- **nsslapd-auditfaillog-compress** for the audit fail log

- **nsslapd-securelog-compress** for the security log

[Bugzilla:1132524](#)

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

Previously, it was difficult to determine who changed an entry if the distinguished name (DN) of the entry did not contain clear identifying information. With the new **nsslapd-auditlog-display-attrs** parameter, you can set additional attributes that Directory Server displays in the audit log to provide 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:

```
time: 20221014125914
dn: uid=73747737483,ou=people,dc=example,dc=com
result: 0
*#cn: John Smith*
changetype: modify
replace: displayName
displayName: jsmith
-
replace: modifiersname
modifiersname: cn=dm
-
replace: modifytimestamp
modifytimestamp: 20221014165914Z
```

Note that if you want the audit log to include all attributes of a modified entry, you can use an asterisk (*) as the parameter value.

[Bugzilla:2136610](#)

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](#)

Directory Server can now import a certificate bundle

Previously, when you tried to add a certificate bundle by using the **dsconf** or **dsctl** utility, the procedure failed with an error, and the certificate bundle was not imported. Such behavior was caused by the **certutil** utility that could import only one certificate at a time. With this update, Directory Server works around the issue with the **certutil**, and a certificate bundle is added successfully.

[Bugzilla:1878808](#)

Default behavior change: Directory Server now returns a DN in exactly the same spelling as it was added to the database

With the new **nsslapd-return-original-entrydn** parameter under the **cn=config** entry, you can manage how Directory Server returns the distinguished name (DN) of entries during search operations.

By default, the **nsslapd-return-original-entrydn** parameter is set to **on**, and Directory Server returns the DN exactly how it was originally added to the database. For example, you added or modified the entry **uid=User,ou=PEople,dc=ExaMPIE,DC=COM**, and with the setting turned on, Directory Server returns the same spelling of the DN for the entry: **uid=User,ou=PEople,dc=ExaMPIE,DC=COM**.

When the **nsslapd-return-original-entrydn** parameter is set to **off**, Directory Server generates the entry DN by putting together a Relative DN (RDN) of the entry and the base DN that is stored in the database suffix configuration under **cn=userroot,cn=ldbm database,cn=plugins,cn=config**. If you set the base DN to **ou=people,dc=example,dc=com**, and the **nsslapd-return-original-entrydn** setting is **off**, Directory Server returns **uid=User,ou=people,dc=example,dc=com** during searches and not the spelling of the DN when you added the entry to the database.

[Bugzilla:2075017](#)

4.14. DESKTOP

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:2154358](#)

Wayland is now enabled on Aspeed GPUs

Previously, the Aspeed GPU driver did not perform well enough to run a Wayland session. To work around that problem, the Wayland session was disabled for Aspeed GPUs.

With this release, the driver performance has been significantly improved and the Wayland session is now responsive. As a result, the Wayland session is now enabled on Aspeed GPUs by default.

[Bugzilla:2131203](#)

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:2160553](#)

4.15. THE WEB CONSOLE

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

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

4.16. RED HAT ENTERPRISE LINUX SYSTEM ROLES

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:2131293](#)

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:2133858](#)

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:2119102](#)

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:2128843](#)

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:2130010](#)

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:2130329](#)

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:2130344](#)

The **logging** RHEL System Role now supports port access and generation of the certificates

With this enhancement, you can use the logging role to manage ports access and generate certificates with new role variables. If you set the new role variables **logging_manage_firewall** and **logging_manage_selinux** to **true**, the roles manage port access. The new role variable for generating certificates is **logging_certificates**. The type and usage are the same as the **certificate** role **certificate_requests**. You can now automate these operations directly by using the **logging** role.

[Bugzilla:2130357](#)

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 manage port access. You can now automate and perform these operations directly by using the **metrics** role.

[Bugzilla:2133528](#)

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:2133930](#)

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:2134202](#)

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:2137663](#)

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:2140795](#)

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:2141330](#)

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:2143768](#)

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:2151282](#)

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:2151283](#)

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:2151284](#)

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:2163709](#)

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:2165175](#)

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:2140804](#)

4.17. 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:1871126](#)

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:1871143](#)

New package: **passt**

This update adds the **passt** package, which makes it possible to use the **passt** user-mode networking back end for virtual machines.

For more information on using **passt**, see the [libvirt upstream documentation](#).

[Bugzilla:2131015](#)

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

4.18. 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](#)

4.19. 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-118705

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-136602

The **container-tools** meta-package has been updated

The **container-tools** RPM meta-package, which contains 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-136607

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-138024

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

Container images now have a two-digit tag

In RHEL 9.0 and RHEL 9.1, container images had a three-digit tag. Starting from RHEL 9.2, container images now have a two-digit tag.

Jira:RHELPLAN-147982

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-129327

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

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-74672

Toolbox can create RHEL 9 containers

Previously, the Toolbox utility only supported RHEL UBI 8 images. With this release, Toolbox now also supports RHEL UBI 9. As a result, you can create a Toolbox container based on RHEL 8 or 9.

The following command creates a RHEL container based on the same RHEL release as your host system:

```
$ toolbox create
```

Alternatively, you can create a container with a specific RHEL release. For example, to create a container based on RHEL 9.2, use the following command:

```
$ toolbox create --distro rhel --release 9.2
```

[Bugzilla:2163752](#)

New package: **passt**

This update adds the **passt** package, which makes it possible to use the **pasta** rootless networking backend for containers.

In comparison to the **Slirp** connection, which is currently used as default for unprivileged networking by Podman, **pasta** provides the following enhancements:

- Improved throughput and better support for IPv6, which includes support for the Neighbor Discovery Protocol (NDP) and for DHCPv6
- The ability to configure port forwarding of TCP and UDP ports on IPv6

To use **pasta** to connect a Podman container, use the **--network pasta** command-line option.

[Bugzilla:2209419](#)

CHAPTER 5. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS

This chapter provides system administrators with a summary of significant changes in the kernel distributed with Red Hat Enterprise Linux 9.2. 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.

printk.console_no_auto_verbose

With this kernel parameter, you can disable console loglevel raise on oops, panic or lockdep-detected issues (only if lock debug is on). With an exception to setups with low baudrate on serial console, set this parameter to **0** to provide more debug information.

- Format: **<bool>**
- Defaults to **0** (**auto_verbose** is enabled)

rcupdate.rcu_exp_cpu_stall_timeout=[KNL]

With this kernel parameter, you can set timeout for expedited RCU CPU stall warning messages. The value is in milliseconds and the maximum allowed value is 21000 milliseconds.

Note that this value is adjusted to an arch timer tick resolution. Setting this to zero causes the value from **rcupdate.rcu_cpu_stall_timeout** to be used (after conversion from seconds to milliseconds).

rcupdate.rcu_task_stall_info=[KNL]

With this parameter, you can set initial timeout in jiffies for RCU task stall informational messages, which give some indication of the problem for those not patient enough to wait for ten minutes. Informational messages are only printed prior to the stall-warning message for a given grace period. Disable with a value less than or equal to zero.

- Defaults to **10** seconds.
- A change in value does not take effect until the beginning of the next grace period.

rcupdate.rcu_task_stall_info_mult=[KNL]

This parameter is a multiplier for time interval between successive RCU task stall informational messages for a given RCU tasks grace period. This value is clamped to one through ten, inclusive.

It defaults to the value of three, so that the first informational message is printed 10 seconds into the grace period, the second at 40 seconds, the third at 160 seconds, and then the stall warning at 600 seconds would prevent a fourth at 640 seconds.

smp.csd_lock_timeout=[KNL]

With this parameter, you can specify the period of time in milliseconds that **smp_call_function()** and friends will wait for a CPU to release the CSD lock. This is useful when diagnosing bugs involving CPUs disabling interrupts for extended periods of time.

- Defaults to **5,000** milliseconds.
- Setting a value of zero disables this feature.
- This feature may be more efficiently disabled using the **csdlock_debug**- kernel parameter.

srcutree.big_cpu_lim=[KNL]

With this parameter, you can specify the number of CPUs constituting a large system, such that **srcu_struct** structures should immediately allocate an **srcu_node** array.

- Defaults to **128**.
- takes effect only when the low-order four bits of **srcutree.convert_to_big** is equal to **3** (decide at boot).

srcutree.convert_to_big=[KNL]

With this parameter, you can specify under what conditions an SRCU tree **srcu_struct** structure will be converted to big form, that is, with an **rcu_node** tree:

- 0: Never.
- 1: At **init_srcu_struct()** time.
- 2: When **rcutorture** decides to.
- 3: Decide at boot time (default).
- 0x1X: Above plus if high contention.
Either way, the **srcu_node** tree will be sized based on the actual runtime number of CPUs (**nr_cpu_ids**) instead of the compile-time **CONFIG_NR_CPUS**.

srcutree.srcu_max_nodelay=[KNL]

With this parameter, you can specify the number of no-delay instances per jiffy for which the SRCU grace period worker thread will be rescheduled with zero delay. Beyond this limit, worker thread will be rescheduled with a sleep delay of one jiffy.

srcutree.srcu_max_nodelay_phase=[KNL]

With this parameter, you can specify the per-grace-period phase, number of non-sleeping polls of readers. Beyond this limit, grace period worker thread will be rescheduled with a sleep delay of one jiffy, between each rescan of the readers, for a grace period phase.

srcutree.srcu_retry_check_delay=[KNL]

With this parameter, you can specify number of microseconds of non-sleeping delay between each non-sleeping poll of readers.

srcutree.small_contention_lim=[KNL]

With this parameter, you can specify the number of update-side contention events per jiffy will be tolerated before initiating a conversion of an **srcu_struct** structure to big form.

**NOTE**

The value of **srcutree.convert_to_big** must have the 0x10 bit set for contention-based conversions to occur.

Updated kernel parameters

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 parameter will be ignored when **crashkernel=X,high** is not used.

deferred_probe_timeout=[KNL]

With this parameter, you can set a timeout in seconds for deferred probe to give up waiting on dependencies to probe. Only specific dependencies (subsystems or drivers) that have opted in will be ignored.

A timeout of **0** will time out at the end of initcalls. If the time out has not expired, the option will be restarted by each successful driver registration. This option will also dump out devices still on the deferred probe list after retrying.

driver_async_probe=[KNL]

With this parameter, you can list of driver names to be probed asynchronously. * (the asterisk) matches with all driver names.

- If * is specified, the rest of the listed driver names are those that will NOT match the *.
Format: **<driver_name1>,<driver_name2>...**

hugetlb_cma=[HW,CMA]

With this parameter, you can specify the size of a CMA area used for allocation of gigantic hugepages. Or using node format, the size of a CMA area per node.

Format: **nn[KMGTPe] or (node format) <node>:nn[KMGTPe][,<node>:nn[KMGTPe]]**

Reserve a CMA area of given size and allocate gigantic hugepages using the CMA allocator. If enabled, the boot-time allocation of gigantic hugepages is skipped.

hugepages=[HW]

With this parameter, you can specify the number of HugeTLB pages to allocate at boot.

- If this follows `hugepagesz`, it specifies the number of pages of `hugepagesz` to be allocated.
- If this is the first HugeTLB parameter on the command line, it specifies the number of pages to allocate for the default huge page size.
- If using node format, the number of pages to allocate per-node can be specified.
See also [Documentation/admin-guide/mm/hugetlbpage.rst](#).

Format: **<integer> or (node format) <node>:<integer>[,<node>:<integer>]**

hugetlb_free_vmemmap=[KNL]

This parameter requires **CONFIG_HUGETLB_PAGE_OPTIMIZE_VMEMMAP** to be enabled. Allows heavy hugetlb users to free up some more memory (7 * PAGE_SIZE for each 2MB hugetlb page).

- Format: { **[oO][Nn]/Y/y/1 | [oO][Ff]/N/n/0 (default)** }
- **[oO][Nn]/Y/y/1**: enable the feature
- **[oO][Ff]/N/n/0**: disable the feature
Built with **CONFIG_HUGETLB_PAGE_OPTIMIZE_VMEMMAP_DEFAULT_ON=y**,

Defaults to *on*.

**NOTE**

This parameter is not compatible with **memory_hotplug.memmap_on_memory**. If both parameters are enabled, **hugetlb_free_vmemmap** takes precedence over **memory_hotplug.memmap_on_memory**.

ivrs_ioapic=[HW,X86-64]

This parameter provides an override to the IOAPIC-ID <-> DEVICE-ID mapping provided in the IVRS ACPI table.

By default, PCI segment is **0**, and can be omitted. For example,

- to map IOAPIC-ID decimal 10 to PCI device 00:14.0, write the parameter as:

```
ivrs_ioapic[10]=00:14.0
```

- to map IOAPIC-ID decimal 10 to PCI segment 0x1 and PCI device 00:14.0, write the parameter as:

```
ivrs_ioapic[10]=0001:00:14.0
```

ivrs_hpet=[HW,X86-64]

This parameter provides an override to the HPET-ID <-> DEVICE-ID mapping provided in the IVRS ACPI table.

By default, PCI segment is **0**, and can be omitted. For example:

- to map HPET-ID decimal 0 to PCI device 00:14.0, write the parameter as:

```
ivrs_hpet[0]=00:14.0
```

- to map HPET-ID decimal 10 to PCI segment 0x1 and PCI device 00:14.0, write the parameter as:

```
ivrs_ioapic[10]=0001:00:14.0
```

ivrs_acpihid=[HW,X86-64]

This parameter provides an override to the ACPI-HID:UID <-> DEVICE-ID mapping provided in the IVRS ACPI table.

For example, to map *UART-HID:UID AMD0020:0* to PCI segment 0x1 and PCI device ID 00:14.5, write the parameter as:

```
ivrs_acpihid[0001:00:14.5]=AMD0020:0
```

By default, PCI segment is **0**, and can be omitted. For example, for the PCI device 00:14.5 write the parameter as:

```
ivrs_acpihid[00:14.5]=AMD0020:0
```

kvm.eager_page_split=[KVM,X86]

With this parameter, you can control whether or not KVM will try to proactively split all huge pages during dirty logging.

Eager page splitting reduces interruptions to vCPU execution by eliminating the write-protection faults and MMU lock contention that would otherwise be required to split huge pages lazily. VM workloads that rarely perform writes or that write only to a small region of VM memory may benefit from disabling eager page splitting to allow huge pages to still be used for reads.

The behavior of eager page splitting depends on whether **KVM_DIRTY_LOG_INITIALLY_SET** is enabled or disabled.

- If disabled, all huge pages in a memslot will be eagerly split when dirty logging is enabled on that memslot.
- If enabled, eager page splitting will be performed during the **KVM_CLEAR_DIRTY** ioctl, and only for the pages being cleared.
Eager page splitting is only supported when **kvm.tdp_mmu=Y**.

Defaults to **Y** (on).

kvm-arm.mode=[KVM,ARM]

With this parameter, you can select one of KVM/arm64's modes of operation.

- none: Forcefully disable KVM.
- nvhe: Standard nVHE-based mode, without support for protected guests.
- protected: nVHE-based mode with support for guests whose state is kept private from the host.

Defaults to **VHE/nVHE** based on hardware support.

nosmep=[X86,PPC64s]

With this parameter, you can disable SMEP (Supervisor Mode Execution Prevention) even if it is supported by processor.

Format: **pci=option[,option...] [PCI] various_PCI_subsystem_options**

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

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

NOTE

- The first format specifies a PCI bus/device/function address which may change if new hardware is inserted, if motherboard firmware changes, or due to changes caused by other kernel parameters. If the domain is left unspecified, it is taken to be zero. Optionally, a path to a device through multiple device and function addresses can be specified after the base address (this is more robust against renumbering issues).
 - The second format selects devices using IDs from the configuration space which may match multiple devices in the system.
- **earlydump**: dump PCI config space before the kernel changes anything
 - **off**: [X86] do not probe for the PCI bus
 - **bios**: [X86-32] force use of PCI BIOS, do not access the hardware directly. Use this if your machine has a non-standard PCI host bridge.
 - **nobios**: [X86-32] disallow use of PCI BIOS, only direct hardware access methods are allowed. Use this if you experience crashes upon bootup and you suspect they are caused by the BIOS.
 - **conf1**: [X86] Force use of PCI Configuration Access Mechanism 1 (configuration address in IO port 0xCF8, data in IO port 0xCFC, both 32-bit).
 - **conf2**: [X86] Force use of PCI Configuration Access Mechanism 2 (IO port 0xCF8 is an 8-bit port for the function, IO port 0xCFA, also 8-bit, sets bus number. The config space is then accessed through ports 0xC000-0xCFFF).
 - See <http://wiki.osdev.org/PCI> for more info on the configuration access mechanisms.
 - **noaer**: [PCIE] If the PCIEAER kernel configuration parameter is enabled, this kernel boot option can be used to disable the use of PCIE advanced error reporting.
 - **nodomains**: [PCI] Disable support for multiple PCI root domains (aka PCI segments, in ACPI-speak).
 - **nommconf**: [X86] Disable use of MMCONFIG for PCI Configuration
 - **check_enable_amd_mmconf** [X86]: check for and enable properly configured MMIO access to PCI config space on AMD family 10h CPU
 - **nomsi**: [MSI] If the **PCI_MSI** kernel configuration parameter is enabled, this kernel boot option can be used to disable the use of MSI interrupts system-wide.

- `noioapicquirk`: [APIC] Disable all boot interrupt quirks. Safety option to keep boot IRQs enabled. This should never be necessary.
- `ioapicreroute`: [APIC] Enable rerouting of boot IRQs to the primary IO-APIC for bridges that cannot disable boot IRQs. This fixes a source of spurious IRQs when the system masks IRQs.
- `noioapicreroute` [APIC] Disable workaround that uses the boot IRQ equivalent of an IRQ that connects to a chipset where boot IRQs cannot be disabled. The opposite of `ioapicreroute`.
- `biosirq`: [X86-32] Use PCI BIOS calls to get the interrupt routing table. These calls are known to be buggy on several machines and they hang the machine when used, but on other computers it is the only way to get the interrupt routing table. Try this option if the kernel is unable to allocate IRQs or discover secondary PCI buses on your Motherboard.
- `rom`: [X86] Assign address space to expansion ROMs. Use with caution as certain devices share address decoders between ROMs and other resources.
- `norom`: [X86] Do not assign address space to expansion ROMs that do not already have BIOS assigned address ranges.
- `nobar`: [X86] Do not assign address space to the BARs that were not assigned by the BIOS.
- `irqmask=0xMMMM`: [X86] Set a bit mask of IRQs allowed to be assigned automatically to PCI devices. You can make the kernel exclude IRQs of your ISA cards this way.
- `pirqaddr=0xAAAAA`: [X86] Specify the physical address of the PIRQ table (normally generated by the BIOS) if it is outside the **F0000h-100000h** range.
- `lastbus=N`: [X86] Scan all buses thru bus #N. Can be useful if the kernel is unable to find your secondary buses and you want to tell it explicitly which ones they are.
- `assign-busses`: [X86] Always assign all PCI bus numbers ourselves, overriding whatever the firmware may have done.
- `usepirqmask`: [X86] Honor the possible IRQ mask stored in the BIOS \$PIR table. This is needed on some systems with broken BIOSes, notably some HP Pavilion N5400 and Omnibook XE3 notebooks. This will have no effect if ACPI IRQ routing is enabled.
- `noacpi`: [X86] Do not use ACPI for IRQ routing or for PCI scanning.
- `use_crs`: [X86] Use PCI host bridge window information from ACPI. On BIOSes from 2008 or later, this is enabled by default. If you need to use this, please report a bug.
- `nocrs`: [X86] Ignore PCI host bridge windows from ACPI. If you need to use this, please report a bug.
- `use_e820`: [X86] Use E820 reservations to exclude parts of PCI host bridge windows. This is a workaround for BIOS defects in host bridge `_CRS` methods. If you need to use this, please report a bug to linux-pci@vger.kernel.org.
- `no_e820`: [X86] Ignore E820 reservations for PCI host bridge windows. This is the default on modern hardware. If you need to use this, please report a bug to linux-pci@vger.kernel.org.
- `routeirq`: Do IRQ routing for all PCI devices. This is normally done in `pci_enable_device()`, so this option is a temporary workaround for broken drivers that do not call it.
- `skip_isa_align`: [X86] do not align io start addr, so can handle more pci cards

- `early: [X86]` Do not do any early type 1 scanning. This might help on some broken boards which machine check when some devices' config space is read. But various workarounds are disabled and some IOMMU drivers will not work.
- `bfsort`: Sort PCI devices into breadth-first order. This sorting is done to get a device order compatible with older (≤ 2.4) kernels.
- `nobfsort`: Do not sort PCI devices into breadth-first order.
- `pcie_bus_tune_off`: Disable PCIe MPS (Max Payload Size) tuning and use the BIOS-configured MPS defaults.
- `pcie_bus_safe`: Set every device's MPS to the largest value supported by all devices below the root complex.
- `pcie_bus_perf` Set device MPS to the largest allowable MPS based on its parent bus. Also set MRRS (Max Read Request Size) to the largest supported value (no larger than the MPS that the device or bus can support) for best performance.
- `pcie_bus_peer2peer`: Set every device's MPS to 128B, which every device is guaranteed to support. This configuration allows peer-to-peer DMA between any pair of devices, possibly at the cost of reduced performance. This also guarantees that hot-added devices will work.
- `cbiosize=nn[KMG]`: The fixed amount of bus space which is reserved for the CardBus bridge's IO window. The default value is 256 bytes.
- `cbmemsize=nn[KMG]`: The fixed amount of bus space which is reserved for the CardBus bridge's memory window. The default value is 64 megabytes.
- `resource_alignment=`
 - Format: **<order of align>@<pci_dev>[; ...]**
 - Specifies alignment and device to reassign aligned memory resources. How to specify the device is described above. If **<order of align>** is not specified, **PAGE_SIZE** is used as alignment. A PCI-PCI bridge can be specified if resource windows need to be expanded. To specify the alignment for several instances of a device, the PCI vendor, device, subvendor, and subdevice may be specified, for example, **12@pci:8086:9c22:103c:198f** for 4096-byte alignment.
- `ecrc=`: Enable/disable PCIe ECRC (transaction layer end-to-end CRC checking).
 - `bios`: Use BIOS/firmware settings. This is the default.
 - `off`: Turn ECRC off
 - `on`: Turn ECRC on.
- `hpiosize=nn[KMG]`: The fixed amount of bus space which is reserved for hotplug bridge's IO window. Default size is 256 bytes.
- `hpmmiosize=nn[KMG]`: The fixed amount of bus space which is reserved for hotplug bridge's MMIO window. Default size is 2 megabytes.
- `hpmmioprefsize=nn[KMG]`: The fixed amount of bus space which is reserved for hotplug bridge's MMIO_PREF window. Default size is 2 megabytes.

- `hpmemsize=nn[KMG]`: The fixed amount of bus space which is reserved for hotplug bridge's MMIO and MMIO_PREF window. Default size is *2 megabytes*.
- `hpbussize=nn`: The minimum amount of additional bus numbers reserved for buses below a hotplug bridge. Default is *1*.
- `realloc=`: Enable/disable reallocating PCI bridge resources if allocations done by BIOS are too small to accommodate resources required by all child devices.
 - `off`: Turn realloc off
 - `on`: Turn realloc on
- `realloc`: same as `realloc=on`
- `noari`: do not use PCIe ARI.
- `noats`: [PCIe, Intel-IOMMU, AMD-IOMMU] do not use PCIe ATS (and IOMMU device IOTLB).
- `pcie_scan_all`: Scan all possible PCIe devices. Otherwise we only look for one device below a PCIe downstream port.
- `big_root_window`: Try to add a big 64bit memory window to the PCIe root complex on AMD CPUs. Some GFX hardware can resize a BAR to allow access to all VRAM. Adding the window is slightly risky (it may conflict with unreported devices), so this taints the kernel.
- `disable_acs_redir=<pci_dev>[; ...]`: Specify one or more PCI devices (in the format specified above) separated by semicolons. Each device specified will have the PCI ACS redirect capabilities forced off which will allow P2P traffic between devices through bridges without forcing it upstream. Note: this removes isolation between devices and may put more devices in an IOMMU group.
- `force_floating`: [S390] Force usage of floating interrupts.
- `nomio`: [S390] Do not use MIO instructions.
- `norid`: [S390] ignore the RID field and force use of one PCI domain per PCI function

`rcupdate.rcu_cpu_stall_timeout=[KNL]`

Set timeout for RCU CPU stall warning messages. The value is in seconds and the maximum allowed value is 300 seconds.

`rcupdate.rcu_task_stall_timeout=[KNL]`

With this parameter, you can set timeout in jiffies for RCU task stall warning messages. Disable with a value less than or equal to zero.

Defaults to **10** minutes.

A change in value does not take effect until the beginning of the next grace period.

`retbleed=[X86]`

With this 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 perf impact. It also enables STIBP if present. Not suitable on Intel.
 - ibpb,nosmt - Like **ibpb** above but will disable SMT when STIBP is not available. This is the alternative for systems which do not have STIBP.
 - 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.
- Selecting **auto** will choose a mitigation method at run time according to the CPU.

Not specifying this option is equivalent to **retbleed=auto**.

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

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. Will be rounded up to a 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

kernel.nmi_wd_lpm_factor (PPC only)

This factor represents the percentage added to **watchdog_thresh** when calculating the NMI watchdog timeout during an LPM. The soft lockup timeout is not impacted. Use this factor to apply to the NMI watchdog timeout (only when **nmi_watchdog** is set to 1).

- A value of **0** means no change.
- Defaults to **200**, which means that the NMI watchdog is set to 30s (based on **watchdog_thresh** equal to 10).

net.core.txrehash

With this parameter, you can control default hash rethink behavior on listening socket when the **SO_TXREHASH** option is set to **SOCK_TXREHASH_DEFAULT** (that is, not overridden by **setsockopt**).

- If set to **1** (default), hash rethink is performed on listening socket.
- If set to **0**, hash rethink is not performed.

net.sctp.reconf_enable - BOOLEAN

With this extension, you can enable or disable extension of Stream Reconfiguration functionality specified in RFC6525. This extension provides the ability to "reset" a stream and includes the parameters of **Outgoing/Incoming SSN Reset**, **SSN/TSN Reset** and **Add Outgoing/Incoming Streams**.

- 1: Enable extension.
- 0: Disable extension.
- Defaults to **0**.

net.sctp.intl_enable - BOOLEAN

With this extension, you can enable or disable extension of User Message Interleaving functionality specified in RFC8260. This extension allows the interleaving of user messages sent on different streams. With this feature enabled, I-DATA chunk will replace DATA chunk to carry user messages if also supported by the peer. Note that to use this feature, you must set this option to **1** and also set socket options **SCTP_FRAGMENT_INTERLEAVE** to **2** and **SCTP_INTERLEAVING_SUPPORTED** to **1**.

- 1: Enable extension.
- 0: Disable extension.
- Defaults to **0**.

net.sctp.ecn_enable - BOOLEAN

With this extension, you can control use of Explicit Congestion Notification (ECN) by SCTP. Like in TCP, ECN is used only when both ends of the SCTP connection indicate support for it. This feature is useful in avoiding losses due to congestion by allowing supporting routers to signal congestion before having to drop packets.

- 1: Enable ecn.
- 0: Disable ecn.
- Defaults to **1**.

vm.hugetlb_optimize_vmemmap

This knob is not available when the **memory_hotplug.memmap_on_memory** kernel parameter is configured or the size of *struct page* (a structure defined in **include/linux/mm_types.h**) is not power of two (an unusual system configuration could result in this).

You can enable (set to 1) or disable (set to 0) the feature of optimizing **vmemmap** pages associated with each HugeTLB page.

- If enabled, the **vmemmap** pages of subsequent allocation of HugeTLB pages from buddy allocator will be optimized (7 pages per 2MB HugeTLB page and 4095 pages per 1GB HugeTLB page), whereas already allocated HugeTLB pages will not be optimized. When those optimized HugeTLB pages are freed from the HugeTLB pool to the buddy allocator, the **vmemmap** pages representing that range needs to be remapped again and the **vmemmap** pages discarded earlier need to be relocated again.
- If your use case is that HugeTLB pages are allocated impromptu (for example, never explicitly allocating HugeTLB pages with **nr_hugepages** but only set **nr_overcommit_hugepages**, those overcommitted HugeTLB pages are allocated

impromptu) instead of being pulled from the HugeTLB pool, you should weigh the benefits of memory savings against the more overhead (~2x slower than before) of allocation or freeing HugeTLB pages between the HugeTLB pool and the buddy allocator. Another behavior to note is that if the system is under heavy memory pressure, it could prevent the user from freeing HugeTLB pages from the HugeTLB pool to the buddy allocator since the allocation of **vmemmap** pages could be failed, you have to retry later if your system encounter this situation.

- If disabled, the **vmemmap** pages of subsequent allocation of HugeTLB pages from buddy allocator will not be optimized meaning the extra overhead at allocation time from buddy allocator disappears, whereas already optimized HugeTLB pages will not be affected. If you want to make sure there are no optimized HugeTLB pages, you can set **nr_hugepages** to **0** first and then disable this. Note that writing **0** to **nr_hugepages** will make any *in use* HugeTLB pages become surplus pages. So, those surplus pages are still optimized until they are no longer in use. You will need to wait for those surplus pages to be released before there are no optimized pages in the system.

net.core.rps_default_mask

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

Changed sysctl parameters

kernel.numa_balancing

With this parameter, you can enable, disable, and configure automatic page fault based NUMA memory balancing. Memory is moved automatically to nodes that access it often. The value to set can be the result of ORing the following:

```
= =====  
0 NUMA_BALANCING_DISABLED  
1 NUMA_BALANCING_NORMAL  
2 NUMA_BALANCING_MEMORY_TIERING  
= =====
```

Or **NUMA_BALANCING_NORMAL** to optimize page placement among different NUMA nodes to reduce remote accessing. On NUMA machines, there is a performance penalty if remote memory is accessed by a CPU. When this feature is enabled the kernel samples what task thread is accessing memory by periodically unmapping pages and later trapping a page fault. At the time of the page fault, it is determined if the data being accessed should be migrated to a local memory node.

Or **NUMA_BALANCING_MEMORY_TIERING** to optimize page placement among different types of memory (represented as different NUMA nodes) to place the hot pages in the fast memory. This is implemented based on unmapping and page fault, too.

net.ipv6.route.max_size

This is now deprecated for ipv6 as garbage collection manages cached route entries.

net.sctp.sctp_wmem

This tunable previously was documented as not having any effect. Now, only the first value (**min**) is used, **default** and **max** are ignored.

- min: Minimum size of send buffer that can be used by SCTP sockets. It is guaranteed to each SCTP socket (but not association) even under moderate memory pressure.
- Defaults to **4K**.

CHAPTER 6. DEVICE DRIVERS

6.1. NEW DRIVERS

- ACPI Video Driver (**video**), only in 64-bit ARM architecture
- CXL driver for CXL memory endpoint devices and switches for memory expansion (**cxl_mem**)
- GNSS receiver core (**gnss**)
- GPIO Simulator Module (**gpio-sim**), only in 64-bit ARM architecture
- VirtIO GPIO driver (**gpio-virtio**), only in 64-bit ARM architecture
- NVIDIA Tegra HTE (Hardware Timestamping Engine) driver (**hte-tegra194**), only in 64-bit ARM architecture
- I2C adapter driver for LPI2C bus (**i2c-imx-lpi2c**), only in 64-bit ARM architecture
- Virtio i2c bus driver (**i2c-virtio**), only in 64-bit ARM architecture
- User level driver support for input subsystem (**uinput**), only in 64-bit ARM architecture
- Module which implements common functions that can be used by the nvme host or target drivers (**nvme-common**)
- AMD PMC Driver (**amd-pmc**), only in AMD and Intel 64-bit architectures
- Nvidia sn2201 platform driver (**nvsw-sn2201**), only in AMD and Intel 64-bit architectures
- Serial multi instantiate pseudo device driver (**serial-multi-instantiate**), only in AMD and Intel 64-bit architectures
- Micro Crystal RV8803 RTC driver (**rtc-rv8803**), only in 64-bit ARM architecture and AMD and Intel 64-bit architectures
- NVIDIA Tegra QSPI Controller Driver (**spi-tegra210-quad**), only in 64-bit ARM architecture
- UCSI driver for Cypress CCGx Type-C controller (**ucsi_ccg**), only in 64-bit ARM architecture
- Confidential computing EFI secret area access (**efi_secret**), only in AMD and Intel 64-bit architectures
- TDX Guest Driver (**tdx-guest**), only in AMD and Intel 64-bit architectures
- HPE watchdog driver (**hpwdt**), only in 64-bit ARM architecture
- POWER Architecture Platform Watchdog Driver (**pseries-wdt**), only in IBM Power Systems, Little Endian

Network drivers

- Driver for VXLAN encapsulated traffic (**vxlan**)
- Marvell OcteonTX2 RVU Admin Function Driver (**rvu_af**), only in 64-bit ARM architecture
- Marvell RVU NIC Physical Function Driver (**rvu_nicpf**), only in 64-bit ARM architecture

- Marvell RVU NIC PTP Driver (**otx2_ptp**), only in 64-bit ARM architecture
- Marvell RVU NIC Virtual Function Driver (**rvu_nicvf**), only in 64-bit ARM architecture
- NVIDIA Tegra MGBE driver (**dwmac-tegra**), only in 64-bit ARM architecture
- Serial line CAN interface (**slcan**), only in 64-bit ARM architecture
- Solarflare Siena network driver (**sfc-siena**), only in IBM Power Systems, Little Endian and AMD and Intel 64-bit architectures

Graphics drivers and miscellaneous drivers

- 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
- DRM DisplayPort AUX bus (**drm_dp_aux_bus**), only in 64-bit ARM architecture
- Host1x driver for Tegra products (**host1x**), only in 64-bit ARM architecture
- NVIDIA Tegra DRM driver (**tegra-drm**), only in 64-bit ARM architecture
- 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
- Intel® auxiliary driver for GSC devices (**mei-gsc**), only in AMD and Intel 64-bit architectures

6.2. UPDATED DRIVERS

Storage driver updates

- Driver for Microchip Smart Family Controller (**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 14.2.0.8 (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.
- CSI debug adapter driver (**scsi_debug**) has been updated to version 0191.
- 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).

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 9. 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	2 (bpf() syscall restricted to privileged users, admin can change)
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_BTFF	y
CONFIG_DEBUG_INFO_BTFF_MODULES	y
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	n
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
Bounded loop support	available
ISA extension v2	available
ISA extension v3	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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

Program type	Available helpers
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_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, 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_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_get_func_ip, bpf_get_attach_cookie, bpf_task_pt_regs, bpf_get_branch_snapshot, bpf_find_vma, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_skb_set_timestamp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data, bpf_tcp_raw_gen_syncookie_ipv4, bpf_tcp_raw_gen_syncookie_ipv6, bpf_tcp_raw_check_syncookie_ipv4, bpf_tcp_raw_check_syncookie_ipv6

Program type	Available helpers
--------------	-------------------

Program type	Available helpers
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_realms, 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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_skb_set_timestamp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data, bpf_tcp_raw_gen_syncookie_ipv4, bpf_tcp_raw_gen_syncookie_ipv6, bpf_tcp_raw_check_syncookie_ipv4, bpf_tcp_raw_check_syncookie_ipv6
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_spin_lock, bpf_spin_unlock, 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_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_get_func_ip, bpf_get_attach_cookie, bpf_task_pt_regs, bpf_get_branch_snapshot, bpf_find_vma, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

Program type	Available helpers
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_sk_to_tcp6_sock, bpf_sk_to_tcp_sock, bpf_sk_to_tcp_timewait_sock, bpf_sk_to_tcp_request_sock, bpf_sk_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_sk_to_unix_sock, bpf_loop, bpf_strncmp, bpf_xdp_get_buff_len, bpf_xdp_load_bytes, bpf_xdp_store_bytes, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_sk_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data, bpf_tcp_raw_gen_syncookie_ipv4, bpf_tcp_raw_gen_syncookie_ipv6, bpf_tcp_raw_check_syncookie_ipv4, bpf_tcp_raw_check_syncookie_ipv6

Program type	Available helpers
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_spin_lock, bpf_spin_unlock, 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_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_get_func_ip, bpf_get_attach_cookie, bpf_task_pt_regs, bpf_get_branch_snapshot, bpf_find_vma, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
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_skc_storage_get, bpf_skc_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_sk_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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

Program type	Available helpers
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

Program type	Available helpers
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_get_netns_cookie, 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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

Program type	Available helpers
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_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_loop, bpf_strncmp, bpf_get_retval, bpf_set_retval, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
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_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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

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_spin_lock, bpf_spin_unlock, 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_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_get_func_ip, bpf_task_pt_regs, bpf_get_branch_snapshot, bpf_find_vma, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

Program type	Available helpers
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
lirc_mode2	not supported
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

Program type	Available helpers
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_loop, bpf_strncmp, bpf_get_retval, bpf_set_retval, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

Program type	Available helpers
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_spin_lock, bpf_spin_unlock, 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_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_get_func_ip, bpf_task_pt_regs, bpf_get_branch_snapshot, bpf_find_vma, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
cgroup_sockopt	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_get_netns_cookie, 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_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_loop, bpf_strncmp, bpf_get_retval, bpf_set_retval, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
tracing	not supported
struct_ops	not supported
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_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data
syscall	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_get_socket_cookie, 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_spin_lock, bpf_spin_unlock, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_send_signal, bpf_skb_output, 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_xdp_output, 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_get_task_stack, bpf_d_path, bpf_copy_from_user, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_sock_from_file, bpf_for_each_map_elem, bpf_snprintf, bpf_sys_bpf, bpf_btf_find_by_name_kind, bpf_sys_close, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_get_func_ip, bpf_task_pt_regs, bpf_get_branch_snapshot, bpf_skc_to_unix_sock, bpf_kallsyms_lookup_name, bpf_find_vma, bpf_loop, bpf_strncmp, bpf_xdp_get_buff_len, bpf_copy_from_user_task, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data

Table 7.3. Available map types

Map type	Available
hash	yes
array	yes

Map type	Available
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
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

Map type	Available
struct_ops	yes
ringbuf	yes
inode_storage	yes
task_storage	yes
bloom_filter	yes

CHAPTER 8. BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 9.2 that have a significant impact on users.

8.1. INSTALLER AND IMAGE CREATION

The installer now displays correct total disk space in Custom partitioning with multipath or DDF RAID devices

Previously, when Custom partitioning was selected in Installer on a system with a multipath or DDF RAID device, the total disk space was not reported correctly and member disk devices were listed as available for partitioning.

With this update, the Custom partitioning in Installer reports correct value for total disk space and only allows to use the DDF RAID or multipath device as a whole.

[Bugzilla:2052938](#)

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:2158210](#)

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:1991843](#)

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:2127100](#)

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:2093793](#)

Anaconda now validates LUKS passphrases for the FIPS requirements

Previously, Anaconda did not check if the length of LUKS passphrases satisfies the FIPS requirements, while the underlying tools performed this check. As a consequence, installing in FIPS mode with a passphrase shorter than 8 characters caused the installer to terminate prematurely.

With this update, the installer has been improved to validate and enforce the minimum length for passphrase. As a result, the installer informs if the LUKS passphrase is too short for use in the FIPS mode and prevents the unexpected termination.

[Bugzilla:2163497](#)

8.2. SUBSCRIPTION MANAGEMENT

Subscription manager no longer denies registration and fetching of Red Hat content

Previously, **subscription-manager** operated in container mode when run under OpenShift Container Platform (OCP) because of improved container detection logic in RHEL 9. As a consequence, the system was unable to use the provided subscription credentials and therefore not fetching Red Hat content.

This update fixed the container detection logic so that **subscription-manager** running under OCP does not detect the system (that is the running pod) as a container. As a result, you can now use the provided subscription credentials or even register using your own credentials to fetch Red Hat content from an OpenShift container.

[Bugzilla:2108549](#)

subscription-manager no longer retains nonessential text in the terminal

Starting with RHEL 9.1, **subscription-manager** displays progress information while processing any operation. Previously, for some languages, typically non-Latin, progress messages did not clean up after the operation finished. With this update, all the messages are cleaned up properly when the operation finishes.

If you have disabled the progress messages before, you can re-enable them by entering the following command:

```
# subscription-manager config --rhsm.progress_messages=1
```

[Bugzilla:2136694](#)

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

[Bugzilla:2111251](#)

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

Previously, the **dnf 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 DNF upgrade transaction for a package group or environment.

[Bugzilla:2122626](#)

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

Patch for [BZ#2108969](#) introduced with [RHBA-2022:8295](#) caused a regression where DNF 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 DNF upgrade no longer skips packages that change architecture from or to **noarch**.

[Bugzilla:2124480](#)

Qt message QM files with 3-letter names are now packaged when an RPM package is being built or rebuilt

Previously, the **find-lang.sh** script could not find Qt message QM files (**.qm**) with names consisting of 3 characters. Consequently, these files were not added to an RPM package.

With this update, the issue has been fixed, and the 3-letter Qt message QM files can now be packaged when building or rebuilding an RPM.

[Bugzilla:2144005](#)

8.4. SHELLS AND COMMAND-LINE TOOLS

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:2172589](#)

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:2160748](#)

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:2127416](#)

8.5. SECURITY

USBGuard saves rules even if RuleFile is not defined

Previously, if the **RuleFolder** configuration directive in USBGuard was set but **RuleFile** 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:2155910](#)

python-sqlalchemy rebased to 1.4.45

The **python-sqlalchemy** package has been rebased to version 1.4.45, which provides many bug fixes over version 1.4.37. Most notably, this version contains a fix for a critical memory bug in the cache key generation.

[Bugzilla:2152649](#)

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:2152635](#)

OpenSSL in SECLEVEL=3 now works with PSK cipher suites

Previously, pre-shared key (PSK) cipher suites were not recognized as performing perfect forward secrecy (PFS) key exchange methods. As a consequence, the **ECDHE-PSK** and **DHE-PSK** cipher suites did not work with OpenSSL configured to **SECLEVEL=3**, for example, when the system-wide

cryptographic policy was set to **FUTURE**. The new version of the **openssl** package fixes this problem.

[Bugzilla:2060044](#)

Clevis now correctly skips commented-out devices in **crypttab**

Previously, Clevis tried to unlock commented-out 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-out devices. Now, if an invalid device is commented out, Clevis does not attempt to unlock it, and **clevis-luks-askpass** finishes appropriately. This makes it easier to troubleshoot and reduces unnecessary service runs.

[Bugzilla:2159728](#)

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:2159735](#)

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:2042345](#)

OOM killer no longer terminates **usbguard** prematurely

Previously, the **usbguard.service** file did not contain a definition of the **OOMScoreAdjust** option for the **systemd** service. Consequently, when the system was low on resources, the **usbguard-daemon** process could be terminated before other unprivileged processes. With this update, **usbguard.service** file now includes **OOMScoreAdjust** setting, which prevents the Out-of-Memory (OOM) killer terminate the **usbguard-daemon** process prematurely.

[Bugzilla:2097419](#)

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:2124488](#)

imklog no longer calls **free()** on missing objects

Previously, the **imklog** module called a **free()** function on an already freed object. Consequently, **imklog** could cause a segmentation fault. With this update, the object is no longer freed twice.

[Bugzilla:2157659](#)

fagenrules --load now works correctly

Previously, the **fapolicyd** service did not correctly handle the signal hang up (SIGHUP). Consequently, **fapolicyd** terminated after receiving SIGHUP, and the **fagenrules --load** command did not work correctly. This update contains a fix for the problem. As a result, **fagenrules --load** now works correctly, and rule updates no longer require manual restarts of **fapolicyd**.

[Bugzilla:2070655](#)

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:2120978](#)

Keylime no longer fails attestation of systems that access multiple IMA-measured files

Previously, if a system that runs the Keylime agent accessed multiple files measured by the Integrity Measurement Architecture (IMA) in quick succession, the Keylime verifier incorrectly processed the IMA log additions. As a consequence, the running hash did not match the correct Platform Configuration Register (PCR) state, and the system failed attestation. This update fixes the problem and systems that quickly access multiple measured files no longer fail attestation.

[Bugzilla:2138167](#)

Keylime policy generation script no longer causes a segmentation fault and core dump

The **create_mb_refstate** script generates policies for measured boot attestation in Keylime. Previously, **create_mb_refstate** incorrectly calculated the data length in the **DevicePath** field. As a consequence, the script tried to access invalid memory using the incorrectly calculated length, which resulted in a segmentation fault and core dump.

This update, which has been published in advisory [RHBA-2022:105318-02](#), prevents the segmentation fault when processing the measured boot event log. As a consequence, you can generate a measured boot policy.

[Bugzilla:2140670](#)

TPM certificates no longer cause Keylime registrar to crash

Previously, some certificates in the Keylime TPM certificate store were malformed x509 certificates and caused the Keylime registrar to crash. This update fixes the problem, and Keylime registrar no longer crashes due to malformed certificates.

[Bugzilla:2142009](#)

8.6. NETWORKING

NetworkManager now preserves IP addresses during reapply before acquiring a new DHCP lease

Previously, after changing the connection settings and then using **nmcli device reapply** command, NetworkManager did not preserve the DHCP lease. Consequently, the IP address got removed temporarily. With this fix, NetworkManager preserves the DHCP lease and uses it until the lease expires or the client requests a new one. As a result, when the **nmcli device reapply** command restarts DHCP client, it does not temporarily remove the IP address.

[Bugzilla:2117352](#)

The firewalld service now triggers the ipset deprecation warning only when using direct rules

Previously, the **firewalld** service used the deprecated **ipset** kernel module when it was not necessary. Consequently, RHEL logged the module's deprecation warning which could be misleading because the **ipset** feature of **firewalld** is not deprecated. With this update, **firewalld** only uses the deprecated **ipset** module and logs the warning if the user explicitly uses **ipsets** with the **--direct** option.

[Bugzilla:2122678](#)

The HNV interface now displays the options after reboot

Previously, the **nmcli** utility created a Hybrid Network Virtualization (HNV) bond by using NetworkManager API. Consequently, after a reboot, the HNV bond lost the primary port setting. With this fix, **nmcli** now uses **hcnmgr** to set bonding options for the primary port. The **hcnmgr** utility supports migration of live partitions with Single Root Input/Output Virtualization (SR-IOV) for hybrid networks. As a result, the HNV bond interface displays the **active slave/primary_reselect** option after reboot.

[Bugzilla:2125152](#)

8.7. KERNEL

FADump enabled with Secure Boot works correctly

Previously, when Firmware Assisted Dump (FADump) was enabled in the Secure Boot environment and any of the booting components exceeded the allocated memory region, system reboots caused a GRUB Out of Memory (OOM) state. This update provides a fix in **kexec-tools** so that Secure Boot and FADump work together correctly.

[Bugzilla:2139000](#)

8.8. BOOT LOADER

grubby now passes arguments to a new kernel correctly

When you add a new kernel using the **grubby** tool and do not specify any arguments, or leave the arguments blank, **grubby** will not pass any arguments to the new kernel and **root** will not be set. Using the **--args** and **--copy-default** options ensures new arguments are appended to the default arguments.

[Bugzilla:2127453](#)

8.9. FILE SYSTEMS AND STORAGE

Installer creating LUKSv2 devices with sector size of 512 bytes

Previously, the RHEL installer created LUKSv2 devices with 4096 bytes sectors if the disk had 4096 bytes physical sectors. With this update, installer now creates LUKSv2 devices with sector size of 512 bytes to offer better disk compatibility with different physical sector sizes used together in one LVM volume group even when the LVM physical volumes are encrypted.

[Bugzilla:2103800](#)

supported_speeds sysfs attribute reports correct speed values

Previously, because of an incorrect definition in the **qla2xxx** driver, the **supported_speeds sysfs** attribute for the HBA reported 20 Gb/s speed instead of the expected 64 Gb/s speed. Consequently, if the HBA supported 64 Gb/s link speed, the **supported_speeds sysfs** value was incorrect, which affected the reported speed value.

With this update, the **supported_speeds sysfs** attribute for HBA reports the correct speed values, which are 16 Gb/s, 32 Gb/s, and 64 Gb/s. You can view the speed values by executing the **cat /sys/class/fc_host/host*/supported_speeds** command.

Bugzilla:2069758

The **lpfc** driver is in a valid state during the **D_ID** port swap

Previously, the SAN Boot host, after issuing the NetApp giveback operation, resulted in LVM hung task warnings and stalled I/O. This problem occurred even when alternate paths were available in a DM-Multipath environment due to the fiber channel **D_ID** port swap. As a consequence of the race condition, the **D_ID** port swap resulted in an inconsistent state in the **lpfc** driver, which prevented I/O from being issued.

With this fix, the **lpfc** driver now ensures a valid state when the **D_ID** port swap occurs. As a result, a fiber channel **D_ID** port swap does not cause hung I/O.

Bugzilla:2173947

8.10. 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:1620043

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 CLB. With this fix, if a cluster property is not set **pcs** displays the default value for the property.

Bugzilla:1796827

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:2133546

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:2125344

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:2125337](#)

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:2092950](#)

8.11. COMPILERS AND DEVELOPMENT TOOLS

DNS lookup can now succeed even when some CNAME records are invalid

Previously, the **glibc** DNS stub resolver treated CNAME records with owner names that are not host names as DNS packet errors. Consequently, the DNS query failed because of the DNS packet errors. With this update, the **glibc** stub resolver now skips invalid CNAME records and the corresponding alias information is not extracted. Therefore, DNS lookups can now succeed even if the server response includes a CNAME chain that contains a domain name that is not a host name.

[Bugzilla:2129005](#)

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:2133019](#)

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.

Note that in RHEL 9, it is recommended to build your applications against the **flexiblas-devel** package and link your projects to the FlexiBLAS wrapper library.

[Bugzilla:2115737](#)

The **tzset** function in **glibc** now sets the daylight variable to a non-zero value if there is any DST rule in the TZ data

Previously, the **tzset** function in **glibc** would set the daylight variable to 0 if the last DST transition in the time zone data file did not result in a clock change due to a simultaneous change in the standard time offset. Consequently, when applications use the daylight variable to check if DST was ever active, they do not get the right result and perform incorrect actions based on this information. To fix this, the **tzset** function now sets the daylight variable to a non-zero value if there is any DST rule in the time zone data, regardless of offset. As a result, applications now observe the presence of DST rules regardless of offset changes.

[Bugzilla:2155352](#)

OpenJDK RSAPSSSignature implementation now validates RSA keys before using them

Previously, the RSAPSSSignature implementation in OpenJDK did not fully check if given RSA keys could be used by the SunRSASign provider before attempting to use them, which would result in errors when using custom security providers. The bug is now fixed and, as a result, the RSAPSSSignature implementation now validates RSA keys and allows other providers to handle these keys when it cannot.

[Bugzilla:2188023](#)

The OpenJDK XML signature provider is now functional in FIPS mode

Previously, the OpenJDK XML signature provider was unable to operate in FIPS mode. As a result of enhancements to FIPS mode support the OpenJDK XML signature provider is now enabled in FIPS mode.

[Bugzilla:2186810](#)

OpenJDK in FIPS mode no longer experiences unexpected errors with certain PKCS#11 tokens

Previously, some PKCS#11 tokens were not fully initialized before use by OpenJDK in FIPS mode resulting in unexpected errors. With this upgrade, these errors are now expected and handled by the FIPS support code.

[Bugzilla:2186806](#)

8.12. 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:2127913](#)

The **dscreate** utility now works correctly when it uses a custom path with the **db_dir** parameter

Previously, an instance that used custom directory paths failed to start because the custom directories had a wrong SELinux label. As a consequence, SELinux denied access to these directories and the instance was not created. With this release, **dscreate** utility sets correct SELinux labels for the custom instance directories.

[Bugzilla:1924569](#)

A password change for the Directory Server replication manager account now works correctly

Previously, after a password change, Directory Server did not properly update the password cache for the replication agreement. As a consequence, when you changed the password for the replication manager account, the replication failed. With this update, Directory Server updates the cache properly and, as a result, the replication works as expected.

[Bugzilla:1956987](#)

The IdM client installer no longer specifies the TLS CA configuration in the **ldap.conf** file

Previously, the IdM client installer specified the TLS CA configuration in the **ldap.conf** file. With this update, OpenLDAP uses the default trust store and the IdM client installer does not set up the TLS CA configuration in the **ldap.conf** file.

[Bugzilla:2094673](#)

8.13. THE WEB CONSOLE

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

In RHEL 9.2.0, 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-2023:4346](#) 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:2207498](#)

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:2126959](#)

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:2143816](#)

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:2153030](#)

Enabling implicit files provider to fix **cockpit-session-recording** SSSD configuration

A disabled SSSD implicit files provider caused the **cockpit-session-recording** 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 **cockpit-session-recording** now works as expected.

[Bugzilla:2153043](#)

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](#)

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:2167528](#)

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:2168735](#)

The **selinux** role now manages policy modules idempotently

Previously, the **selinux** role copied an existing module to the managed node every time, reporting a change even when the module was already present. With this update, the **selinux** role checks if the module has been installed on the managed node, and does not attempt to copy and install the module if it is already installed.

[Bugzilla:2160152](#)

The **rhc** system role no longer fails on the registered systems when **rhc_auth** contains activation keys

Previously, a failure occurred when you executed playbook files on the registered systems with the activation key specified in the **rhc_auth** parameter. This issue has been resolved. It is now possible to execute playbook files on the already registered systems, even when activation keys are provided in the **rhc_auth** parameter.

[Bugzilla:2186218](#)

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:2140899](#)

VMs on IBM Z no longer fail to start when using **memfd** memory backing

Previously, on IBM Z hosts, virtual machines (VMs) failed to boot if they were configured to use the **memfd** type of hugepage memory backing, for example as follows:

```
<memoryBacking>
  <hugepages/>
  <source type='memfd'/>
</memoryBacking>
```

With this update, the underlying cause has been fixed, and the affected VMs now start correctly.

[Bugzilla:2116496](#)

VNC can now reliably connect to UEFI VMs after migration

Previously, if you enabled or disabled a message queue while migrating a virtual machine (VM), the Virtual Network Computing (VNC) client failed to connect to the VM after the migration was complete.

This problem affected only UEFI-based VMs that used the Open Virtual Machine Firmware (OVMF).

The problem has been fixed, and the VNC client now reliably connects to UEFI VMs after the migration is complete.

Jira:RHELPLAN-135600

The installer shows the expected system disk to install RHEL on VM

Previously, when installing RHEL on a VM using **virtio-scsi** devices, it was possible that these devices did not appear in the installer because of a **device-mapper-multipath** bug. Consequently, during installation, if some devices had a serial set and some did not, the **multipath** command was claiming all the devices that had a serial. Due to this, the installer was unable to find the expected system disk to install RHEL in the VM.

With this update, **multipath** correctly sets the devices with no serial as having no World Wide Identifier (WWID) and ignores them. On installation, **multipath** only claims devices that **multipathd** uses to bind a multipath device, and the installer shows the expected system disk to install RHEL in the VM.

Bugzilla:1926147

CHAPTER 9. TECHNOLOGY PREVIEWS

This part provides a list of all Technology Previews available in Red Hat Enterprise Linux 9.

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

9.1. INSTALLER AND IMAGE CREATION

NVMe over Fibre Channel devices are now available in RHEL installer as a Technology Preview

You can now add NVMe over Fibre Channel devices to your RHEL installation as a Technology Preview. In RHEL Installer, you can select these devices under the NVMe Fabrics Devices section while adding disks on the Installation Destination screen.

[Bugzilla:2107346](#)

9.2. SHELLS AND COMMAND-LINE TOOLS

GIMP available as a Technology Preview in RHEL 9

GNU Image Manipulation Program (GIMP) 2.99.8 is now available in RHEL 9 as a Technology Preview. The **gimp** package version 2.99.8 is a pre-release version with a set of improvements, but a limited set of features and no guarantee for stability. As soon as the official GIMP 3 is released, it will be introduced into RHEL 9 as an update of this pre-release version.

In RHEL 9, you can install **gimp** easily as an RPM package.

[Bugzilla:2047161](#)

9.3. 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](#)

9.4. SECURITY

gnutls now uses KTLS as a Technology Preview

The updated **gnutls** packages can use Kernel TLS (KTLS) for accelerating data transfer on encrypted channels as a Technology Preview. To enable KTLS, add the **tls.ko** kernel module using the **modprobe** command, and create a new configuration file **/etc/crypto-policies/local.d/gnutls-ktls.txt** for the system-wide cryptographic policies with the following content:

```
[global]
ktls = true
```

Note that the current version does not support updating traffic keys through TLS **KeyUpdate** messages, which impacts the security of AES-GCM ciphersuites. See the [RFC 7841 - TLS 1.3](#) document for more information.

Bugzilla:2042009

9.5. NETWORKING

WireGuard VPN is available as a Technology Preview

WireGuard, which Red Hat provides as an unsupported Technology Preview, is a high-performance VPN solution that runs in the Linux kernel. It uses modern cryptography and is easier to configure than other VPN solutions. Additionally, the small code-basis of WireGuard reduces the surface for attacks and, therefore, improves the security.

For further details, see [Setting up a WireGuard VPN](#).

Bugzilla:1613522

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

The **systemd-resolved** service is 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 **systemd-resolved** is an unsupported Technology Preview.

[Bugzilla:2020529](#)

9.6. KERNEL

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. The version 1 enables platforms using the **Flexible Launch Control** mechanism to use the SGX technology.

Bugzilla:1874182

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 the shared work queue with process address space ID (pasid) submission and shared virtual memory (SVM).

[Bugzilla:2030412](#)

The Soft-iWARP driver is available as a Technology Preview

Soft-iWARP (siw) is a software, Internet Wide-area RDMA Protocol (iWARP), kernel driver for Linux. Soft-iWARP implements the iWARP protocol suite over the TCP/IP network stack. This protocol suite is fully implemented in software and does not require a specific Remote Direct Memory Access (RDMA) hardware. Soft-iWARP enables a system with a standard Ethernet adapter to connect to an iWARP adapter or to another system with already installed Soft-iWARP.

Bugzilla:2023416

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

rvu_af, rvu_nicpf, and rvu_nicvf available as Technology Preview

The following kernel modules are available as Technology Preview for Marvell OCTEON TX2 Infrastructure Processor family:

- **rvu_nicpf** - Marvell OcteonTX2 NIC Physical Function driver
- **rvu_nicvf** - Marvell OcteonTX2 NIC Virtual Function driver
- **rvu_nicvf** - Marvell OcteonTX2 RVU Admin Function driver

Bugzilla:2040643

9.7. FILE SYSTEMS AND STORAGE

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

In RHEL 9, the DAX file system is available as a Technology Preview. DAX provides 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 DAX compatible file system must be created on the NVDIMM(s). Also, the file system must be mounted with the **dax** mount option. Then, an **mmap** of a file on the dax-mounted file system results in a direct mapping of storage into the application's address space.

Bugzilla:1995338

Stratis is available as a Technology Preview

Stratis is a local storage manager. It provides managed file systems on top of pools of storage with additional features to the user:

- Manage snapshots and thin provisioning

- Automatically grow file system sizes as needed
- Maintain file systems

To administer Stratis storage, use the **stratis** utility, which communicates with the **stratisd** background service.

Stratis is provided as a Technology Preview.

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

[Bugzilla:2041558](#)

NVMe-oF Discovery Service features available as a Technology Preview

The NVMe-oF Discovery Service features, defined in the NVMexpress.org Technical Proposals (TP) 8013 and 8014, are available as a Technology Preview. To preview these features, use the **nvme-cli 2.0** package and attach the host to an NVMe-oF target device that implements TP-8013 or TP-8014. For more information about TP-8013 and TP-8014, see the NVM Express 2.0 Ratified TPs from the <https://nvmexpress.org/specifications/> website.

[Bugzilla:2021672](#)

nvme-stas package available as a Technology Preview

The **nvme-stas** package, which is a Central Discovery Controller (CDC) client for Linux, is now available as a Technology Preview. It handles Asynchronous Event Notifications (AEN), Automated NVMe subsystem connection controls, Error handling and reporting, and Automatic (**zeroconf**) and Manual configuration.

This package consists of two daemons, Storage Appliance Finder (**stafd**) and Storage Appliance Connector (**stacd**).

[Bugzilla:1893841](#)

NVMe TP 8006 in-band authentication available as a Technology Preview

Implementing Non-Volatile Memory Express (NVMe) TP 8006, which is an in-band authentication for NVMe over Fabrics (NVMe-oF) is now available as an unsupported Technology Preview. The NVMe Technical Proposal 8006 defines the **DH-HMAC-CHAP** in-band authentication protocol for NVMe-oF, which is provided with this enhancement.

For more information, see the **dhchap-secret** and **dhchap-ctrl-secret** option descriptions in the **nvme-connect(1)** man page.

[Bugzilla:2027304](#)

9.8. COMPILERS AND DEVELOPMENT TOOLS

jmc-core and owasp-java-encoder available as a Technology Preview

RHEL 9 is distributed with the **jmc-core** and **owasp-java-encoder** packages as Technology Preview features for the AMD and Intel 64-bit architectures.

jmc-core is a library providing core APIs for Java Development Kit (JDK) Mission Control, including libraries for parsing and writing JDK Flight Recording files, as well as libraries for Java Virtual Machine (JVM) discovery through Java Discovery Protocol (JDP).

The **owasp-java-encoder** package provides a collection of high-performance low-overhead contextual encoders for Java.

Note that since RHEL 9.2, **jmc-core** and **owasp-java-encoder** are available in the CodeReady Linux Builder (CRB) repository, which you must explicitly enable. See [How to enable and make use of content within CodeReady Linux Builder](#) for more information.

[Bugzilla:1980981](#)

9.9. IDENTITY MANAGEMENT

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:2084180](#)

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:2084166](#)

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 9.1 and later.

[Bugzilla:2065693](#)

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 9.1 and later.

[Bugzilla:2056482](#)

RHEL IdM allows delegating user authentication to external identity providers 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 9.1 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:2069202](#)

9.10. 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](#)

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

9.11. GRAPHICS INFRASTRUCTURES

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:2041690

9.12. THE WEB CONSOLE

Stratis available as a Technology Preview in the RHEL web console

With this update, the Red Hat Enterprise Linux web console provides the ability to manage Stratis storage as a Technology Preview.

To learn more about Stratis, see [What is Stratis](#).

Jira:RHELPLAN-122345

9.13. VIRTUALIZATION

Intel SGX available for VMs as a Technology Preview

As a Technology Preview, the Intel Software Guard Extensions (SGX) can now be configured for virtual machines (VMs) hosted on RHEL 9. SGX helps protect data integrity and confidentiality for specific processes on Intel hardware. After you set up SGX on your host, the feature is passed on to its VMs, so that the guest operating systems (OSs) can use it.

Note that for a guest OS to use SGX, you must first install SGX drivers for that specific OS. In addition, SGX on your host cannot memory-encrypt VMs.

Jira:RHELPLAN-69761

AMD SEV and SEV-ES for KVM virtual machines

As a Technology Preview, RHEL 9 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 9 includes SEV and SEV-ES encryption, but not the SEV and SEV-ES security attestation.

Jira:RHELPLAN-65217

Virtualization is now available on ARM 64

As a Technology Preview, it is now possible to create KVM virtual machines on systems using ARM 64 CPUs.

Jira:RHELPLAN-103993

virtio-mem is now available on AMD64, Intel 64, and ARM 64

As a Technology Preview, RHEL 9 introduces the **virtio-mem** feature on AMD64, Intel 64, and ARM 64 systems. Using **virtio-mem** makes it possible to dynamically add or remove host memory in virtual machines (VMs).

To use **virtio-mem**, define **virtio-mem** memory devices in the XML configuration of a VM and use the **virsh update-memory-device** command to request memory device size changes while the VM is running. To see the current memory size exposed by such memory devices to a running VM, view the XML configuration of the VM.

[Bugzilla:2014487](#), [Bugzilla:2044172](#), [Bugzilla:2044162](#)

Intel TDX in RHEL guests

As a Technology Preview, the Intel Trust Domain Extension (TDX) feature can now be used in RHEL 9.2 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:1955275

A unified kernel image of RHEL is now available as a Technology Preview

As a Technology Preview, you can now obtain the RHEL kernel as a unified kernel image (UKI) for virtual machines (VMs). A unified kernel image combines the kernel, initramfs, and kernel command line into a single signed binary file.

UKIs can be used in virtualized and cloud environments, especially in confidential VMs where strong SecureBoot capabilities are required. The UKI is available as a **kernel-uki-virt** package in RHEL 9 repositories.

Currently, the RHEL UKI can only be used in a UEFI boot configuration.

Bugzilla:2142102

9.14. RHEL IN CLOUD ENVIRONMENTS

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

With the updated RHEL kernel, you can now create and run RHEL confidential virtual machines (VMs) on Microsoft Azure as a Technology Preview. The newly added unified kernel image (UKI) now enables booting encrypted confidential VM images on Azure. The UKI is available as a **kernel-uki-virt** package in RHEL 9 repositories.

Currently, the RHEL UKI can only be used in a UEFI boot configuration.

Jira:RHELPLAN-139800

9.15. CONTAINERS

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

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-136611

CHAPTER 10. DEPRECATED FUNCTIONALITY

This part provides an overview of functionality that has been *deprecated* in Red Hat Enterprise Linux 9.

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 9. 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 8 but has been *removed* in RHEL 9, see [Considerations in adopting RHEL 9](#).

10.1. INSTALLER AND IMAGE CREATION

Deprecated Kickstart commands

The following Kickstart commands have been deprecated:

- **timezone --ntpservers**
- **timezone --nntp**
- **logging --level**
- **%packages --excludeWeakdeps**
- **%packages --instLangs**
- **%anaconda**
- **pwpolicy**

Note that where only specific options are listed, the base command and its other options are still available and not deprecated. Using the deprecated commands in Kickstart files prints a warning in the logs. You can turn the deprecated command warnings into errors with the **inst.ksstrict** boot option.

Bugzilla:1899167

User and Group customizations in the **edge-commit** and **edge-container** blueprints have been deprecated

Specifying a user or group customization in the blueprints is deprecated for the **edge-commit** and **edge-container** image types, because the user customization disappears when you upgrade the image and do not specify the user in the blueprint again. Therefore, you should specify the users and groups

directly in the blueprints for edge image types which are used to deploy an existing OSTree commit, such as **edge-raw-image**, **edge-installer**, and **edge-simplified-installer**.

Note that specifying a user or group customization in blueprints remains supported, but the support will be eventually removed.

[Bugzilla:2173928](#)

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:2163716](#)

10.3. SHELLS AND COMMAND-LINE TOOLS

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 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:2089395](#)

10.4. SECURITY

SHA-1 is deprecated for cryptographic purposes

The usage of the SHA-1 message digest for cryptographic purposes has been deprecated in RHEL 9.

The digest produced by SHA-1 is not considered secure because of many documented successful attacks based on finding hash collisions. The RHEL core crypto components no longer create signatures using SHA-1 by default. Applications in RHEL 9 have been updated to avoid using SHA-1 in security-relevant use cases.

Among the exceptions, the HMAC-SHA1 message authentication code and the Universal Unique Identifier (UUID) values can still be created using SHA-1 because these use cases do not currently pose security risks. SHA-1 also can be used in limited cases connected with important interoperability and compatibility concerns, such as Kerberos and WPA-2. See the [List of RHEL applications using cryptography that is not compliant with FIPS 140-3](#) section in the [RHEL 9 Security hardening document](#) for more details.

If your scenario requires the use of SHA-1 for verifying existing or third-party cryptographic signatures, you can enable it by entering the following command:

```
# update-crypto-policies --set DEFAULT:SHA1
```

Alternatively, you can switch the system-wide crypto policies to the **LEGACY** policy. Note that **LEGACY** also enables many other algorithms that are not secure.

Jira:RHELPLAN-110763

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:2054740](#)

SCP is deprecated in RHEL 9

The secure copy protocol (SCP) is deprecated because it has known security vulnerabilities. The SCP API remains available for the RHEL 9 lifecycle but using it reduces system security.

- In the **scp** utility, SCP is replaced by the SSH File Transfer Protocol (SFTP) by default.
- The OpenSSH suite does not use SCP in RHEL 9.
- SCP is deprecated in the **libssh** library.

Jira:RHELPLAN-99136

Digest-MD5 in SASL is deprecated

The Digest-MD5 authentication mechanism in the Simple Authentication Security Layer (SASL) framework is deprecated, and it might be removed from the **cyrus-sasl** packages in a future major release.

[Bugzilla:1995600](#)

OpenSSL deprecates MD2, MD4, MDC2, Whirlpool, Blowfish, CAST, DES, IDEA, RC2, RC4, RC5, SEED, and PBKDF1

The OpenSSL project has deprecated a set of cryptographic algorithms because they are insecure, uncommonly used, or both. Red Hat also discourages the use of those algorithms, and RHEL 9 provides them for migrating encrypted data to use new algorithms. Users must not depend on those algorithms

for the security of their systems.

The implementations of the following algorithms have been moved to the legacy provider in OpenSSL: MD2, MD4, MDC2, Whirlpool, Blowfish, CAST, DES, IDEA, RC2, RC4, RC5, SEED, and PBKDF1.

See the **/etc/pki/tls/openssl.cnf** configuration file for instructions on how to load the legacy provider and enable support for the deprecated algorithms.

[Bugzilla:1975836](#)

/etc/system-fips is now deprecated

Support for indicating FIPS mode through the **/etc/system-fips** file has been removed, and the file will not be included in future versions of RHEL. To install RHEL in FIPS mode, add the **fips=1** parameter to the kernel command line during the system installation. You can check whether RHEL operates in FIPS mode by using the **fips-mode-setup --check** command.

Jira:RHELPLAN-103232

libcrypt.so.1 is now deprecated

The **libcrypt.so.1** library is now deprecated, and it might be removed in a future version of RHEL.

[Bugzilla:2034569](#)

OpenSSL requires padding for RSA encryption in FIPS mode

OpenSSL no longer supports RSA encryption without padding in FIPS mode. RSA encryption without padding is uncommon and is rarely used. Note that key encapsulation with RSA (RSASVE) does not use padding but is still supported.

[Bugzilla:2168665](#)

10.5. NETWORKING

Network teams are deprecated in RHEL 9

The **teamd** service and the **libteam** library are deprecated in Red Hat Enterprise Linux 9 and will be removed in the next major release. As a replacement, configure a bond instead of a network team.

Red Hat focuses its efforts on kernel-based bonding to avoid maintaining two features, bonds and teams, that have similar functions. The bonding code has a high customer adoption, is robust, and has an active community development. As a result, the bonding code receives enhancements and updates.

For details about how to migrate a team to a bond, see [Migrating a network team configuration to network bond](#).

[Bugzilla:1935544](#)

NetworkManager connection profiles in ifcfg format are deprecated

In RHEL 9.0 and later, connection profiles in **ifcfg** format are deprecated. The next major RHEL release will remove the support for this format. However, in RHEL 9, NetworkManager still processes and updates existing profiles in this format if you modify them.

By default, NetworkManager now stores connection profiles in keyfile format in the **/etc/NetworkManager/system-connections/** directory. Unlike the **ifcfg** format, the keyfile format supports all connection settings that NetworkManager provides. For further details about the keyfile

format and how to migrate profiles, see [NetworkManager connection profiles in keyfile format](#) .

Bugzilla:1894877

The iptables back end in firewalld is deprecated

In RHEL 9, the **iptables** framework is deprecated. As a consequence, the **iptables** backend and the **direct interface** in **firewalld** are also deprecated. Instead of the **direct interface** you can use the native features in **firewalld** to configure the required rules.

[Bugzilla:2089200](#)

10.6. KERNEL

ATM encapsulation is deprecated in RHEL 9

Asynchronous Transfer Mode (ATM) encapsulation enables Layer-2 (Point-to-Point Protocol, Ethernet) or Layer-3 (IP) connectivity for the ATM Adaptation Layer 5 (AAL-5). Red Hat has not been providing support for ATM NIC drivers since RHEL 7. The support for ATM implementation is being dropped in RHEL 9. These protocols are currently used only in chipsets, which support the ADSL technology and are being phased out by manufacturers. Therefore, ATM encapsulation is deprecated in Red Hat Enterprise Linux 9.

For more information, see [PPP Over AAL5, Multiprotocol Encapsulation over ATM Adaptation Layer 5](#) , and [Classical IP and ARP over ATM](#) .

[Bugzilla:2058153](#)

The kexec_load system call for kexec-tools has been deprecated

The **kexec_load** system call, which loads the second kernel, will not be supported in future RHEL releases. The **kexec_file_load** system call replaces **kexec_load** and is now the default system call on all architectures.

Bugzilla:2113873

Network teams are deprecated in RHEL 9

The **teamd** service and the **libteam** library are deprecated in Red Hat Enterprise Linux 9 and will be removed in the next major release. As a replacement, configure a bond instead of a network team.

Red Hat focuses its efforts on kernel-based bonding to avoid maintaining two features, bonds and teams, that have similar functions. The bonding code has a high customer adoption, is robust, and has an active community development. As a result, the bonding code receives enhancements and updates.

For details about how to migrate a team to a bond, see [Migrating a network team configuration to network bond](#).

Bugzilla:2013884

10.7. FILE SYSTEMS AND STORAGE

lvm2-activation-generator and its generated services removed in RHEL 9.0

The **lvm2-activation-generator** program and its generated services **lvm2-activation**, **lvm2-activation-early**, and **lvm2-activation-net** are removed in RHEL 9.0. The **lvm.conf event_activation** setting, used to activate the services, is no longer functional. The only method for auto activating volume groups is

event based activation.

[Bugzilla:2038183](#)

10.8. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

libdb has been deprecated

RHEL 8 and RHEL 9 currently provide Berkeley DB (**libdb**) version 5.3.28, which is distributed under the LGPLv2 license. The upstream Berkeley DB version 6 is available under the AGPLv3 license, which is more restrictive.

The **libdb** package is deprecated as of RHEL 9 and might not be available in future major RHEL releases.

In addition, cryptographic algorithms have been removed from **libdb** in RHEL 9 and multiple **libdb** dependencies have been removed from RHEL 9.

Users of **libdb** are advised to migrate to a different key-value database. For more information, see the Knowledgebase article [Available replacements for the deprecated Berkeley DB \(libdb\) in RHEL](#) .

[Bugzilla:1927780](#), [Jira:RHELPLAN-80695](#), [Bugzilla:1974657](#)

10.9. COMPILERS AND DEVELOPMENT TOOLS

Smaller size of keys than 2048 are deprecated by openssl 3.0

Key sizes smaller than 2048 bits are deprecated by **openssl** 3.0 and no longer work in Go's FIPS mode.

[Bugzilla:2111072](#)

Some PKCS1 v1.5 modes are now deprecated

Some **PKCS1** v1.5 modes are not approved in **FIPS-140-3** for encryption and are disabled. They will no longer work in Go's FIPS mode.

[Bugzilla:2092016](#)

10.10. IDENTITY MANAGEMENT

SHA-1 in OpenDNSSec is now deprecated

OpenDNSSec supports exporting Digital Signatures and authentication records using the **SHA-1** algorithm. The use of the **SHA-1** algorithm is no longer supported. With the RHEL 9 release, **SHA-1** in OpenDNSSec is deprecated and it might be removed in a future minor release. Additionally, OpenDNSSec support is limited to its integration with Red Hat Identity Management. OpenDNSSec is not supported standalone.

[Bugzilla:1979521](#)

The SSSD implicit files provider domain is disabled by default

The SSSD implicit **files** provider domain, which retrieves user information from local files such as **/etc/shadow** and group information from **/etc/groups**, is now disabled by default.

To retrieve user and group information from local files with SSSD:

1. Configure SSSD. Choose one of the following options:

- a. Explicitly configure a local domain with the **id_provider=files** option in the **sssd.conf** configuration file.

```
[domain/local]
id_provider=files
...
```

- b. Enable the **files** provider by setting **enable_files_domain=true** in the **sssd.conf** configuration file.

```
[sssd]
enable_files_domain = true
```

2. Configure the name services switch.

```
# authselect enable-feature with-files-provider
```

Jira:RHELPLAN-100639

-h and -p options were deprecated in OpenLDAP client utilities.

The upstream OpenLDAP project has deprecated the **-h** and **-p** options in its utilities, and recommends using the **-H** option instead to specify the LDAP URI. As a consequence, RHEL 9 has deprecated these two options in all OpenLDAP client utilities. The **-h** and **-p** options will be removed from RHEL products in future releases.

Jira:RHELPLAN-137660

The SSSD files provider has been deprecated

The SSSD **files** provider has been deprecated in Red Hat Enterprise Linux (RHEL) 9. The **files** provider might be removed from a future release of RHEL.

Jira:RHELPLAN-139805

The nsslapd-idlistscanlimit parameter is deprecated and its default value has been changed

With the new filter reordering optimization, the **nsslapd-idlistscanlimit** attribute impact on search performance is more harmful than helpful. As a result, the attribute is deprecated. Additionally, the default value has been changed to **2147483646** (unlimited).

[Bugzilla:1952241](#)

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:RHELDPCS-16612

10.11. DESKTOP

GTK 2 is now deprecated

The legacy GTK 2 toolkit and the following, related packages have been deprecated:

- **adwaita-gtk2-theme**
- **gnome-common**
- **gtk2**
- **gtk2-immodules**
- **hexchat**

Several other packages currently depend on GTK 2. These have been modified so that they no longer depend on the deprecated packages in a future major RHEL release.

If you maintain an application that uses GTK 2, Red Hat recommends that you port the application to GTK 4.

Jira:RHELPLAN-131882

LibreOffice is deprecated

The LibreOffice RPM packages are now deprecated and will be removed in a future major RHEL release. LibreOffice continues to be fully supported through the entire life cycle of RHEL 7, 8, and 9.

As a replacement for the RPM packages, Red Hat recommends that you install LibreOffice from either of the following sources provided by The Document Foundation:

- The official Flatpak package in the Flathub repository:
<https://flathub.org/apps/org.libreoffice.LibreOffice>.
- The official RPM packages: <https://www.libreoffice.org/download/download-libreoffice/>.

Jira:RHELDPCS-16300

10.12. GRAPHICS INFRASTRUCTURES

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

10.13. RED HAT ENTERPRISE LINUX SYSTEM ROLES

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 a RHEL 8 control node to configure a network team on RHEL 9 nodes, shows a warning about the deprecation.

[Bugzilla:1999770](#)

10.14. VIRTUALIZATION

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](#)

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. However, a new VM snapshot mechanism is under development and is planned to be fully implemented in a future minor release of RHEL 9.

Jira:RHELPLAN-15509, [Bugzilla:1621944](#)

The virtual floppy driver has become deprecated

The **isa-fdc** driver, which controls virtual floppy disk devices, is now deprecated, and will become unsupported in a future release of RHEL. Therefore, to ensure forward compatibility with migrated virtual machines (VMs), Red Hat discourages using floppy disk devices in VMs hosted on RHEL 9.

[Bugzilla:1965079](#)

qcow2-v2 image format is deprecated

With RHEL 9, the qcow2-v2 format for virtual disk images has become deprecated, and will become unsupported in a future major release of RHEL. In addition, the RHEL 9 Image Builder cannot create disk images in the qcow2-v2 format.

Instead of qcow2-v2, Red Hat strongly recommends using qcow2-v3. To convert a qcow2-v2 image to a later format version, use the **qemu-img amend** command.

[Bugzilla:1951814](#)

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

libvirt has become deprecated

The monolithic **libvirt** daemon, **libvirtd**, has been deprecated in RHEL 9, and will be removed in a future major release of RHEL. Note that you can still use **libvirtd** for managing virtualization on your hypervisor, but Red Hat recommends switching to the newly introduced modular **libvirt** daemons. For instructions and details, see the [RHEL 9 Configuring and Managing Virtualization](#) document.

Jira:RHELPLAN-113995

Legacy CPU models are now deprecated

A significant number of CPU models have become deprecated and will become unsupported for use in virtual machines (VMs) in a future major release of RHEL. The deprecated models are as follows:

- For Intel: models prior to Intel Xeon 55xx and 75xx Processor families (also known as Nehalem)
- For AMD: models prior to AMD Opteron G4
- For IBM Z: models prior to IBM z14

To check whether your VM is using a deprecated CPU model, use the **virsh dominfo** utility, and look for a line similar to the following in the **Messages** section:

```
tainted: use of deprecated configuration settings
deprecated configuration: CPU model 'i486'
```

[Bugzilla:2060839](#)

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

10.15. CONTAINERS

Running RHEL 9 containers on a RHEL 7 host is not supported

Running RHEL 9 containers on a RHEL 7 host is not supported. It might work, but it is not guaranteed.

For more information, see [Red Hat Enterprise Linux Container Compatibility Matrix](#) .

Jira:RHELPLAN-100087

SHA1 hash algorithm within Podman has been deprecated

The SHA1 algorithm used to generate the filename of the rootless network namespace is no longer supported in Podman. Therefore, rootless containers started before updating to Podman 4.1.1 or later have to be restarted if they are joined to a network (and not just using **slirp4netns**) to ensure they can connect to containers started after the upgrade.

Bugzilla:2069279

rhel9/pause has been deprecated

The **rhel9/pause** container image has been deprecated.

[Bugzilla:2106816](#)

The CNI network stack has been deprecated

The Container Network Interface (CNI) network stack has been deprecated. 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-147725

10.16. 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 8 and RHEL 9, see [Changes to packages](#) in the *Considerations in adopting RHEL 9* document.



IMPORTANT

The support status of deprecated packages remains unchanged within RHEL 9. 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 9:

- iptables-devel
- iptables-libs
- iptables-nft
- iptables-nft-services
- iptables-utils
- libdb
- mcpp
- mod_auth_mellon

- `motif`
- `motif-devel`
- `python3-pytz`
- `xorg-x11-server-Xorg`

CHAPTER 11. KNOWN ISSUES

This part describes known issues in Red Hat Enterprise Linux 9.2.

11.1. INSTALLER AND IMAGE CREATION

The **auth** and **authconfig** Kickstart commands require the AppStream repository

The **authselect-compat** package is required by the **auth** and **authconfig** Kickstart commands during installation. Without this package, the installation fails if **auth** or **authconfig** are used. However, by design, the **authselect-compat** package is only available in the AppStream repository.

To work around this problem, verify that the BaseOS and AppStream repositories are available to the installer or use the **authselect** Kickstart command during installation.

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

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](#)

Local Media installation source is not detected when booting the installation from a USB that is created using a third party tool

When booting the RHEL installation from a USB that is created using a third party tool, the installer fails to detect the **Local Media** installation source (only *Red Hat CDN* is detected).

This issue occurs because the default boot option **int.stage2=** attempts to search for **iso9660** image format. However, a third party tool might create an ISO image with a different format.

As a workaround, use either of the following solution:

- When booting the installation, click the **Tab** key to edit the kernel command line, and change the boot option **inst.stage2=** to **inst.repo=**.
- To create a bootable USB device on Windows, use Fedora Media Writer.

- When using a third party tool like Rufus to create a bootable USB device, first regenerate the RHEL ISO image on a Linux system, and then use the third party tool to create a bootable USB device.

For more information on the steps involved in performing any of the specified workaround, see, [Installation media is not auto detected during the installation of RHEL 8.3](#) .

Bugzilla:1877697

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

Driver disk menu fails to display user inputs on the console

When you start RHEL installation using the **inst.dd** option on the Kernel command line with a driver disk, the console fails to display the user input. Consequently, it appears that the application does not respond to the user input and freezes, but displays the output which is confusing for users. However, this behavior does not affect the functionality, and user input gets registered after pressing **Enter**.

As a workaround, to see the expected results, ignore the absence of user inputs in the console and press **Enter** when you finish adding inputs.

Bugzilla:2109231

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 workaround 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

Anaconda fails to verify existence of an administrator user account

While installing RHEL using a graphical user interface, Anaconda fails to verify if the administrator account has been created. As a consequence, users might install a system without any administrator user account.

To work around this problem, ensure you configure an administrator user account or the root password is set and the root account is unlocked. As a result, users can perform administrative tasks on the installed system.

[Bugzilla:2047713](#)

New XFS features prevent booting of PowerNV IBM POWER systems with firmware older than version 5.10

PowerNV IBM POWER systems use a Linux kernel for firmware, and use Petitboot as a replacement for GRUB. This results in the firmware kernel mounting **/boot** and Petitboot reading the GRUB config and booting RHEL.

The RHEL 9 kernel introduces **bigtime=1** and **inobtcount=1** features to the XFS filesystem, which kernels with firmware older than version 5.10 do not understand.

To work around this problem, you can use another filesystem for **/boot**, for example ext4.

[Bugzilla:1997832](#)

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:2125542](#)

NetworkManager fails to start after the installation when connected to a network but without DHCP or a static IP address configured

Starting with RHEL 9.0, Anaconda activates network devices automatically when there is no specific **ip=** or kickstart network configuration set. Anaconda creates a default persistent configuration file for each Ethernet device. The connection profile has the **ONBOOT** and **autoconnect** value set to **true**. As a consequence, during the start of the installed system, RHEL activates the network devices, and the **networkManager-wait-online** service fails.

As a workaround, do one of the following:

- Delete all connections using the **nmcli** utility except one connection you want to use. For example:
 - a. List all connection profiles:

```
# nmcli connection show
```


- b. Delete the connection profiles that you do not require:

```
# nmcli connection delete <connection_name>
```

Replace `<connection_name>` with the name of the connection you want to delete.

- Disable the auto connect network feature in Anaconda if no specific **ip=** or kickstart network configuration is set.
 - a. In the Anaconda GUI, navigate to **Network & Host Name**
 - b. Select a network device to disable.
 - c. Click **Configure**.
 - d. On the **General** tab, deselect the **Connect automatically with priority**
 - e. Click **Save**.

Bugzilla:2115783

Unable to load an updated driver from the driver update disc in the installation environment

A new version of a driver from the driver update disc might not load if the same driver from the installation initial ramdisk has already been loaded. As a consequence, an updated version of the driver cannot be applied to the installation environment.

As a workaround, use the **modprobe.blacklist=** kernel command line option together with the **inst.dd** option. For example, to ensure that an updated version of the **virtio_blk** driver from a driver update disc is loaded, use **modprobe.blacklist=virtio_blk** and then continue with the usual procedure to apply drivers from the driver update disk. As a result, the system can load an updated version of the driver and use it in the installation environment.

Bugzilla:2164216

Kickstart installations fail to configure the network connection

Anaconda performs the kickstart network configuration only through the NetworkManager API. Anaconda processes the network configuration after the **%pre** kickstart section. As a consequence, some tasks from the kickstart **%pre** section are blocked. For example, downloading packages from the **%pre** section fails due to unavailability of the network configuration.

To work around this problem:

- Configure the network, for example using the **nmcli** tool, as a part of the **%pre** script.
- Use the installer boot options to configure the network for the **%pre** script.

As a result, it is possible to use the network for tasks in the **%pre** section and the kickstart installation process completes.

Bugzilla:2173992

11.2. SOFTWARE MANAGEMENT

The Installation process sometimes becomes unresponsive

When you install RHEL, the installation process sometimes becomes unresponsive. The `/tmp/packaging.log` file displays the following message at the end:

```
10:20:56,416 DDEBUG dnf: RPM transaction over.
```

To workaroud this problem, restart the installation process.

[Bugzilla:2073510](#)

11.3. SHELLS AND COMMAND-LINE TOOLS

Renaming network interfaces using `ifcfg` files fails

On RHEL 9, the **initscripts** package is not installed by default. Consequently, renaming network interfaces using `ifcfg` files fails. To solve this problem, Red Hat recommends that you use **udev** rules or link files to rename interfaces. For further details, see [Consistent network interface device naming](#) and the **systemd.link(5)** man page.

If you cannot use one of the recommended solutions, install the **initscripts** package.

[Bugzilla:2018112](#)

The **chkconfig** package is not installed by default in RHEL 9

The **chkconfig** package, which updates and queries runlevel information for system services, is not installed by default in RHEL 9.

To manage services, use the **systemctl** commands or install the **chkconfig** package manually.

For more information about **systemd**, see [Introduction to systemd](#). For instructions on how to use the **systemctl** utility, see [Managing system services with systemctl](#).

[Bugzilla:2053598](#)

The Service Location Protocol (SLP) is vulnerable to an attack through UDP

The OpenSLP provides a dynamic configuration mechanism for applications in local area networks, such as printers and file servers. However, SLP is vulnerable to a reflective denial of service amplification attack through UDP on systems connected to the internet. SLP allows an unauthenticated attacker to register new services without limits set by the SLP implementation. By using UDP and spoofing the source address, an attacker can request the service list, creating a Denial of Service on the spoofed address.

To prevent external attackers from accessing the SLP service, disable SLP on all systems running on untrusted networks, such as those directly connected to the internet. Alternatively, to work around this problem, configure firewalls to block or filter traffic on UDP and TCP port 427.

[Bugzilla:2184570](#)

11.4. INFRASTRUCTURE SERVICES

Both **bind** and **unbound** disable validation of SHA-1-based signatures

The **bind** and **unbound** components disable validation support of all RSA/SHA1 (algorithm number 5) and RSASHA1-NSEC3-SHA1 (algorithm number 7) signatures, and the SHA-1 usage for signatures is restricted in the DEFAULT system-wide cryptographic policy.

As a result, certain DNSSEC records signed with the SHA-1, RSA/SHA1, and RSASHA1-NSEC3-SHA1 digest algorithms fail to verify in Red Hat Enterprise Linux 9 and the affected domain names become vulnerable.

To work around this problem, upgrade to a different signature algorithm, such as RSA/SHA-256 or elliptic curve keys.

For more information and a list of top-level domains that are affected and vulnerable, see the [DNSSEC records signed with RSASHA1 fail to verify](#) solution.

[Bugzilla:2070495](#)

named fails to start if the same writable zone file is used in multiple zones

BIND does not allow the same writable zone file in multiple zones. Consequently, if a configuration includes multiple zones which share a path to a file that can be modified by the **named** service, **named** fails to start. To work around this problem, use the **in-view** clause to share one zone between multiple views and make sure to use different paths for different zones. For example, include the view names in the path.

Note that writable zone files are typically used in zones with allowed dynamic updates, slave zones, or zones maintained by DNSSEC.

[Bugzilla:1984982](#)

libotr is not compliant with FIPS

The **libotr** library and toolkit for off-the-record (OTR) messaging provides end-to-end encryption for instant messaging conversations. However, the **libotr** library does not conform to the Federal Information Processing Standards (FIPS) due to its use of the **gcry_pk_sign()** and **gcry_pk_verify()** functions. As a result, you cannot use the **libotr** library in FIPS mode.

[Bugzilla:2086562](#)

Setting the console keymap requires the libxkbcommon library on your minimal install

In RHEL 9, certain **systemd** library dependencies have been converted from dynamic linking to dynamic loading, so that your system opens and uses the libraries at runtime when they are available. With this change, a functionality that depends on such libraries is not available unless you install the necessary library. This also affects setting the keyboard layout on systems with a minimal install. As a result, the **localectl --no-convert set-x11-keymap gb** command fails.

To work around this problem, install the **libxkbcommon** library:

```
# dnf install libxkbcommon
```

[Bugzilla:2214130](#)

The %vmeff metric from the sysstat package displays incorrect values

The **sysstat** package provides the **%vmeff** metric to measure the page reclaim efficiency. The values of the **%vmeff** column returned by the **sar -B** command are incorrect because **sysstat** does not parse all relevant **/proc/vmstat** values provided by later kernel versions. To work around this problem, you can calculate the **%vmeff** value manually from the **/proc/vmstat** file. For details, see [Why the sar\(1\) tool reports %vmeff values beyond 100 % in RHEL 8 and RHEL 9?](#)

[Bugzilla:2230431](#)

11.5. 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](#)

OpenSSL does not detect if a PKCS #11 token supports the creation of raw RSA or RSA-PSS signatures

The TLS 1.3 protocol requires support for RSA-PSS signatures. If a PKCS #11 token does not support raw RSA or RSA-PSS signatures, server applications that use the OpenSSL library fail to work with an RSA key if the key is held by the PKCS #11 token. As a result, TLS communication fails in the described scenario.

To work around this problem, configure servers and clients to use TLS version 1.2 as the highest TLS protocol version available.

[Bugzilla:1681178](#)

OpenSSL incorrectly handles PKCS #11 tokens that does not support raw RSA or RSA-PSS signatures

The **OpenSSL** library does not detect key-related capabilities of PKCS #11 tokens. Consequently, establishing a TLS connection fails when a signature is created with a token that does not support raw RSA or RSA-PSS signatures.

To work around the problem, add the following lines after the **.include** line at the end of the **crypto_policy** section in the **/etc/pki/tls/openssl.cnf** file:

```
SignatureAlgorithms =  
RSA+SHA256:RSA+SHA512:RSA+SHA384:ECDSA+SHA256:ECDSA+SHA512:ECDSA+SHA384  
MaxProtocol = TLSv1.2
```

As a result, a TLS connection can be established in the described scenario.

[Bugzilla:1685470](#)

scp empties files copied to themselves when a specific syntax is used

The **scp** utility changed from the Secure copy protocol (SCP) to the more secure SSH file transfer protocol (SFTP). Consequently, copying a file from a location to the same location erases the file content. The problem affects the following syntax:

scp localhost:/myfile localhost:/myfile

To work around this problem, do not copy files to a destination that is the same as the source location using this syntax.

The problem has been fixed for the following syntaxes:

- `scp /myfile localhost:/myfile`
- `scp localhost:~/myfile ~/myfile`

[Bugzilla:2056884](#)

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:2165920](#)

Ansible remediations require additional collections

With the replacement of Ansible Engine by the **ansible-core** package, the list of Ansible modules provided with the RHEL subscription is reduced. As a consequence, running remediations that use Ansible content included within the **scap-security-guide** package requires collections from the **rhc-worker-playbook** package.

For an Ansible remediation, perform the following steps:

1. Install the required packages:

```
# dnf install -y ansible-core scap-security-guide rhc-worker-playbook
```

2. Navigate to the `/usr/share/scap-security-guide/ansible` directory:

```
# cd /usr/share/scap-security-guide/ansible
```

3. Run the relevant Ansible playbook using environment variables that define the path to the additional Ansible collections:

```
# ANSIBLE_COLLECTIONS_PATH=/usr/share/rhc-worker-
  playbook/ansible/collections/ansible_collections/ ansible-playbook -c local -i localhost, rhel9-
  playbook-cis_server_11.yml
```

Replace **cis_server_11** with the ID of the profile against which you want to remediate the system.

As a result, the Ansible content is processed correctly.

**NOTE**

Support of the collections provided in **rhc-worker-playbook** is limited to enabling the Ansible content sourced in **scap-security-guide**.

[Bugzilla:2105162](#)

oscap-anaconda-addon does not allow CIS hardening of systems with Network Servers package group

When installing RHEL Network Servers with a CIS security profile (**cis**, **cis_server_I1**, **cis_workstation_I1**, or **cis_workstation_I2**) on systems with the Network Servers package group selected, **oscap-anaconda-addon** sends the error message **package tftp has been added to the list of excluded packages, but it can't be removed from the current software selection without breaking the install**. To proceed with the installation, navigate back to Software Selection and uncheck the **Network Servers** additional software to allow the installation and hardening to finish. Then, install the required packages.

[Bugzilla:2172264](#)

Keylime does not accept concatenated PEM certificates

When Keylime receives a certificate chain as multiple certificates in the PEM format concatenated in a single file, the **keylime-agent-rust** Keylime component does not correctly use all the provided certificates during signature verification, resulting in a TLS handshake failure. As a consequence, the client components (**keylime_verifier** and **keylime_tenant**) cannot connect to the Keylime agent. To work around this problem, use just one certificate instead of multiple certificates.

Jira:RHELPLAN-157225

Keylime requires a specific file for `tls_dir = default`

When the **tls_dir** variable is set to **default** in Keylime verifier or registrar configuration, Keylime checks for the presence of the **cacert.crt** file in the **/var/lib/keylime/cv_ca** directory. If the file is not present, the **keylime_verifier** or **keylime_registrar** service fails to start and records the following message in a log: **Exception: It appears that the verifier has not yet created a CA and certificates, please run the verifier first**. As a consequence, Keylime rejects custom certificate authority (CA) certificates that have a different file name even when they are placed in the **/var/lib/keylime/ca_cv** directory.

To work around this problem and use custom CA certificates, manually specify **tls_dir = /var/lib/keylime/ca_cv** instead of using **tls_dir = default**.

Jira:RHELPLAN-157337

Default SELinux policy allows unconfined executables to make their stack executable

The default state of the **selinuxuser_execstack** boolean in the SELinux policy is on, which means that unconfined executables can make their stack executable. Executables should not use this option, and it might indicate poorly coded executables or a possible attack. However, due to compatibility with other tools, packages, and third-party products, Red Hat cannot change the value of the boolean in the default policy. If your scenario does not depend on such compatibility aspects, you can turn the boolean off in your local policy by entering the command **setsebool -P selinuxuser_execstack off**.

[Bugzilla:2064274](#)

SSH timeout rules in STIG profiles configure incorrect options

An update of OpenSSH affected the rules in the following Defense Information Systems Agency Security Technical Implementation Guide (DISA STIG) profiles:

- DISA STIG for RHEL 9 (**xccdf_org.ssgproject.content_profile_stig**)
- DISA STIG with GUI for RHEL 9 (**xccdf_org.ssgproject.content_profile_stig_gui**)

In each of these profiles, the following two rules are affected:

Title: Set SSH Client Alive Count Max to zero
 CCE Identifier: CCE-90271-8
 Rule ID: **xccdf_org.ssgproject.content_rule_sshd_set_keepalive_0**

Title: Set SSH Idle Timeout Interval
 CCE Identifier: CCE-90811-1
 Rule ID: **xccdf_org.ssgproject.content_rule_sshd_set_idle_timeout**

When applied to SSH servers, each of these rules configures an option (**ClientAliveCountMax** and **ClientAliveInterval**) that no longer behaves as previously. As a consequence, OpenSSH no longer disconnects idle SSH users when it reaches the timeout configured by these rules. As a workaround, these rules have been temporarily removed from the DISA STIG for RHEL 9 and DISA STIG with GUI for RHEL 9 profiles until a solution is developed.

[Bugzilla:2038978](#)

GnuPG incorrectly allows using SHA-1 signatures even if disallowed by crypto-policies

The GNU Privacy Guard (GnuPG) cryptographic software can create and verify signatures that use the SHA-1 algorithm regardless of the settings defined by the system-wide cryptographic policies. Consequently, you can use SHA-1 for cryptographic purposes in the **DEFAULT** cryptographic policy, which is not consistent with the system-wide deprecation of this insecure algorithm for signatures.

To work around this problem, do not use GnuPG options that involve SHA-1. As a result, you will prevent GnuPG from lowering the default system security by using the non-secure SHA-1 signatures.

[Bugzilla:2070722](#)

gpg-agent does not work as an SSH agent in FIPS mode

The **gpg-agent** tool creates MD5 fingerprints when adding keys to the **ssh-agent** program even though FIPS mode disables the MD5 digest. Consequently, the **ssh-add** utility fails to add the keys to the authentication agent.

To work around the problem, create the **~/.gnupg/sshcontrol** file without using the **gpg-agent --daemon --enable-ssh-support** command. For example, you can paste the output of the **gpg --list-keys** command in the **<FINGERPRINT> 0** format to **~/.gnupg/sshcontrol**. As a result, **gpg-agent** works as an SSH authentication agent.

[Bugzilla:2073567](#)

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](#)

11.6. NETWORKING

The **nm-cloud-setup** service removes manually-configured secondary IP addresses from interfaces

Based on the information received from the cloud environment, the **nm-cloud-setup** service configures network interfaces. Disable **nm-cloud-setup** to manually configure interfaces. However, in certain cases, other services on the host can configure interfaces as well. For example, these services could add secondary IP addresses. To avoid that **nm-cloud-setup** removes secondary IP addresses:

1. Stop and disable the **nm-cloud-setup** service and timer:

```
# systemctl disable --now nm-cloud-setup.service nm-cloud-setup.timer
```

2. Display the available connection profiles:

```
# nmcli connection show
```

3. Reactive the affected connection profiles:

```
# nmcli connection up "<profile_name>"
```

As a result, the service no longer removes manually-configured secondary IP addresses from interfaces.

[Bugzilla:2151040](#)

Failure to update the session key causes the connection to break

Kernel Transport Layer Security (kTLS) protocol does not support updating the session key, which is used by the symmetric cipher. Consequently, the user cannot update the key, which causes a connection break. To work around this problem, disable kTLS. As a result, with the workaround, it is possible to successfully update the session key.

[Bugzilla:2013650](#)

The **initscripts** package is not installed by default

By default, the **initscripts** package is not installed. As a consequence, the **ifup** and **ifdown** utilities are not available. As an alternative, use the **nmcli connection up** and **nmcli connection down** commands to enable and disable connections. If the suggested alternative does not work for you, report the problem and install the **NetworkManager-initscripts-updown** package, which provides a NetworkManager solution for the **ifup** and **ifdown** utilities.

[Bugzilla:2082303](#)

11.7. KERNEL

The **kdump** mechanism in kernel causes OOM errors on the 64K kernel

The 64K kernel page size on the 64-bit ARM architecture uses more memory than the 4KB kernel. Consequently, **kdump** causes a kernel panic and memory allocation fails with out of memory (OOM) errors. As a work around, manually configure the **crashkernel** value to 640 MB. For example, set the **crashkernel=** parameter as **crashkernel=2G- :640M**.

As a result, the **kdump** mechanism does not fail on the 64K kernel in the described scenario.

[Bugzilla:2160676](#)

Customer applications with dependencies on kernel page size may need updating when moving from 4k to 64k page size kernel

RHEL is compatible with both 4k and 64k page size kernels. Customer applications with dependencies on a 4k kernel page size may require updating when moving from 4k to 64k page size kernels. Known instances of this include **jemalloc** and dependent applications.

The **jemalloc** memory allocator library is sensitive to the page size used in the system's runtime environment. The library can be built to be compatible with 4k and 64k page size kernels, for example, when configured with **--with-lg-page=16** or **env JEMALLOC_SYS_WITH_LG_PAGE=16** (for **jemallocator** Rust crate). Consequently, a mismatch can occur between the page size of the runtime environment and the page size that was present when compiling binaries that depend on **jemalloc**. As a result, using a **jemalloc**-based application triggers the following error:

```
<jemalloc>: Unsupported system page size
```

To avoid this problem, use one of the following approaches:

- Use the appropriate build configuration or environment options to create 4k and 64k page size compatible binaries.
- Build any userspace packages that use **jemalloc** after booting into the final 64k kernel and runtime environment.

For example, you can build the **fd-find** tool, which also uses **jemalloc**, with the **cargo** Rust package manager. In the final 64k environment, trigger a new build of all dependencies to resolve the mismatch in the page size by entering the **cargo** command:

```
# cargo install fd-find --force
```

[Bugzilla:2167783](#)

The **kdump** service fails to build the **initrd** file on IBM Z systems

On the 64-bit IBM Z systems, the **kdump** service fails to load the initial RAM disk (**initrd**) when **znet** related configuration information such as **s390-subchannels** reside in an inactive **NetworkManager** connection profile. Consequently, the **kdump** mechanism fails with the following error:

```
dracut: Failed to set up znet
kdump: mkdumprd: failed to make kdump initrd
```

As a workaround, use one of the following solutions:

- Configure a network bond or bridge by re-using the connection profile that has the **znet** configuration information:

```
$ nmcli connection modify enc600 master bond0 slave-type bond
```

- Copy the **znet** configuration information from the inactive connection profile to the active connection profile:
 - a. Run the **nmcli** command to query the **NetworkManager** connection profiles:

```
# nmcli connection show

NAME                UUID                TYPE  Device
bridge-br0          ed391a43-bdea-4170-b8a2 bridge  br0
bridge-slave-enc600 caf7f770-1e55-4126-a2f4 ethernet enc600
enc600              bc293b8d-ef1e-45f6-bad1 ethernet --
```

- b. Update the active profile with configuration information from the inactive connection:

```
#!/bin/bash
inactive_connection=enc600
active_connection=bridge-slave-enc600
for name in nettype subchannels options; do
field=802-3-ethernet.s390-$name
val=$(nmcli --get-values "$field"connection show "$inactive_connection")
nmcli connection modify "$active_connection" "$field" $val
done
```

- c. Restart the **kdump** service for changes to take effect:

```
# kdumpctl restart
```

[Bugzilla:2064708](#)

kTLS does not support offloading of TLS 1.3 to NICs

Kernel Transport Layer Security (kTLS) does not support offloading of TLS 1.3 to NICs. Consequently, software encryption is used with TLS 1.3 even when the NICs support TLS offload. To work around this problem, disable TLS 1.3 if offload is required. As a result, you can offload only TLS 1.2. When TLS 1.3 is in use, there is lower performance, since TLS 1.3 cannot be offloaded.

[Bugzilla:2000616](#)

The Delay Accounting functionality does not display the SWAPIN and IO% statistics columns by default

The **Delayed Accounting** functionality, unlike early versions, is disabled by default. Consequently, the **iotop** application does not show the **SWAPIN** and **IO%** statistics columns and displays the following warning:

```
CONFIG_TASK_DELAY_ACCT not enabled in kernel, cannot determine SWAPIN and IO%
```

The **Delay Accounting** functionality, using the **taskstats** interface, provides the delay statistics for all tasks or threads that belong to a thread group. Delays in task execution occur when they wait for a kernel resource to become available, for example, a task waiting for a free CPU to run on. The statistics help in setting a task's CPU priority, I/O priority, and **rss** limit values appropriately.

As a workaround, you can enable the **delayacct** boot option either at run time or boot.

- To enable **delayacct** at run time, enter:

```
echo 1 > /proc/sys/kernel/task_delayacct
```

Note that this command enables the feature system wide, but only for the tasks that you start after running this command.

- To enable **delayacct** permanently at boot, use one of the following procedures:
 - Edit the **/etc/sysctl.conf** file to override the default parameters:
 - a. Add the following entry to the **/etc/sysctl.conf** file:

```
kernel.task_delayacct = 1
```

For more information, see [How to set sysctl variables on Red Hat Enterprise Linux](#) .

- b. Reboot the system for changes to take effect.
 - Add the **delayacct** option to the kernel command line.

For more information, see [Configuring kernel command-line parameters](#).

As a result, the **iotop** application displays the **SWAPIN** and **IO%** statistics columns.

Bugzilla:2132480

The **kdump** mechanism fails to capture the **vmcore** file on LUKS-encrypted targets

When running **kdump** on systems with Linux Unified Key Setup (LUKS) encrypted partitions, systems require a certain amount of available memory. When the available memory is less than the required amount of memory, the **systemd-cryptsetup** service fails to mount the partition. Consequently, the second kernel fails to capture the crash dump file (**vmcore**) on LUKS-encrypted targets.

With the **kdumpctl estimate** command, you can query the **Recommended crashkernel value**, which is the recommended memory size required for **kdump**.

To work around this problem, use following steps to configure the required memory for **kdump** on LUKS encrypted targets:

1. Print the estimate **crashkernel** value:

```
# kdumpctl estimate
```

2. Configure the amount of required memory by increasing the **crashkernel** value:

```
# grubby --args=crashkernel=652M --update-kernel=ALL
```

3. Reboot the system for changes to take effect.

```
# reboot
```

As a result, **kdump** works correctly on systems with LUKS-encrypted partitions.

Bugzilla:2017401

Allocating crash kernel memory fails at boot time

On certain Ampere Altra systems, allocating the crash kernel memory for **kdump** usage fails during boot when the available memory is below 1 GB. Consequently, the **kdumpctl** command fails to start the **kdump** service.

To workaround this problem, do one of the following:

- Decrease the value of the **crashkernel** parameter by a minimum of 240 MB to fit the size requirement, for example **crashkernel=240M**.
- Use the **crashkernel=x,high** option to reserve crash kernel memory above 4 GB for **kdump**.

As a result, the crash kernel memory allocation for **kdump** does not fail on Ampere Altra systems.

[Bugzilla:2065013](#)

RHEL fails to recognize NVMe disks when VMD is enabled

When you reset or reattach the driver, the Volume Management Device (VMD) domain currently does not soft-reset. Consequently, the hardware cannot properly detect and enumerate its devices. As a result, the operating system with VMD enabled does not recognize NVMe disks, especially when resetting a server or working with a VM machine.

Bugzilla:2128610

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 9.1 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:2129288

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

The **mlx5** driver fails while using the Mellanox **ConnectX-5** adapter

In Ethernet switch device driver model (**switchdev**) mode, the **mlx5** driver fails when configured with the device managed flow steering (DMFS) parameter and **ConnectX-5** adapter supported hardware. As a consequence, you can see the following error message:

```
BUG: Bad page cache in process umount pfn:142b4b
```

To work around this problem, use the software managed flow steering (SMFS) parameter instead of DMFS.

Bugzilla:2180665

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

11.8. BOOT LOADER

Cannot install RHEL when PReP is not 4 or 8 MiB in size

The RHEL installer cannot install the boot loader if the PowerPC Reference Platform (PReP) partition is of a different size than 4 MiB or 8 MiB on a disk that uses 4 kiB sectors. As a consequence, you cannot install RHEL on the disk.

To work around the problem, make sure that the PReP partition is exactly 4 MiB or 8 MiB in size, and that the size is not rounded to another value. As a result, the installer can now install RHEL on the disk.

Bugzilla:2026579

11.9. FILE SYSTEMS AND STORAGE

Anaconda fails to login iSCSI server using the **no authentication** method after unsuccessful CHAP authentication attempt

When you add iSCSI discs using CHAP authentication and the login attempt fails due to incorrect credentials, a relogin attempt to the discs with the **no authentication** method fails. To work around this problem, close the current session and login using the **no authentication** method.

Bugzilla:1983602

Device Mapper Multipath is not supported with NVMe/TCP

Using Device Mapper Multipath with the **nvme-tcp** driver can result in the Call Trace warnings and system instability. To work around this problem, NVMe/TCP users must enable native NVMe multipathing and not use the **device-mapper-multipath** tools with NVMe.

By default, Native NVMe multipathing is enabled in RHEL 9. For more information, see [Enabling multipathing on NVMe devices](#).

Bugzilla:2033080

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

Disabling quota accounting is no longer possible for an XFS filesystem mounted with quotas enabled

As of RHEL 9.2, it is no longer possible to disable quota accounting on an XFS filesystem which has been mounted with quotas enabled.

To work around this issue, disable quota accounting by remounting the filesystem, with the quota option removed.

Bugzilla:2160619

System fails to boot when adding an NVMe-FC device as a mount point in **/etc/fstab**

The Non-volatile Memory Express over Fibre Channel (NVMe-FC) devices mounted through the **/etc/fstab** file fails to mount at boot and the system enters into emergency mode. This is due to a known bug in the **nvme-cli nvmf-autoconnect systemd** services.

Bugzilla:2168603

udev rule change for NVMe devices

There is a udev rule change for NVMe devices that adds **OPTIONS="string_escape=replace"** parameter. This leads to a disk by-id naming change for some vendors, if the serial number of your device has leading whitespace.

[Bugzilla:2185048](#)

11.10. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

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:2157708](#)

The --ssl-fips-mode option in MySQL and MariaDB does not change FIPS mode

The **--ssl-fips-mode** option in **MySQL** and **MariaDB** in RHEL works differently than in upstream.

In RHEL 9, if you use **--ssl-fips-mode** as an argument for the **mysqld** or **mariadb** daemon, or if you use **ssl-fips-mode** in the **MySQL** or **MariaDB** server configuration files, **--ssl-fips-mode** does not change FIPS mode for these database servers.

Instead:

- If you set **--ssl-fips-mode** to **ON**, the **mysqld** or **mariadb** server daemon does not start.
- If you set **--ssl-fips-mode** to **OFF** on a FIPS-enabled system, the **mysqld** or **mariadb** server daemons still run in FIPS mode.

This is expected because FIPS mode should be enabled or disabled for the whole RHEL system, not for specific components.

Therefore, do not use the **--ssl-fips-mode** option in **MySQL** or **MariaDB** in RHEL. Instead, ensure FIPS mode is enabled on the whole RHEL system:

- Preferably, install RHEL with FIPS mode enabled. Enabling FIPS mode during the installation ensures that the system generates all keys with FIPS-approved algorithms and continuous monitoring tests in place. For information about installing RHEL in FIPS mode, see [Installing the system in FIPS mode](#).
- Alternatively, you can switch FIPS mode for the entire RHEL system by following the procedure in [Switching the system to FIPS mode](#).

[Bugzilla:1991500](#)

11.11. COMPILERS AND DEVELOPMENT TOOLS

Certain symbol-based probes do not work in **SystemTap** on the 64-bit ARM architecture

Kernel configuration disables certain functionality needed for **SystemTap**. Consequently, some symbol-based probes do not work on the 64-bit ARM architecture. As a result, affected **SystemTap** scripts may not run or may not collect hits on desired probe points.

Note that this bug has been fixed for the remaining architectures with the release of the [RHBA-2022:5259](#) advisory.

Bugzilla:2083727

GCC in GCC Toolset 12: CPU detection may fail on Intel Sapphire Rapids processors

CPU detection on Intel Sapphire Rapids processors relies on the existence of the **AVX512_VP2INTERSECT** feature. This feature has been removed from the GCC Toolset 12 version of GCC and, as a consequence, CPU detection may fail on Intel Sapphire Rapids processors.

[Bugzilla:2141718](#)

11.12. IDENTITY MANAGEMENT

Configuring a referral for a suffix fails in Directory Server

If you set a back-end referral in Directory Server, setting the state of the backend using the **dsconf <instance_name> backend suffix set --state referral** command fails with the following error:

```
Error: 103 - 9 - 53 - Server is unwilling to perform - [] - need to set nsslapd-referral before moving to referral state
```

As a consequence, configuring a referral for suffixes fail. To work around the problem:

1. Set the **nsslapd-referral** parameter manually:

```
# ldapmodify -D "cn=Directory Manager" -W -H ldap://server.example.com

dn: cn=dc\3Dexample\2Cdc\3Dcom,cn=mapping tree,cn=config
changetype: modify
add: nsslapd-referral
nsslapd-referral: ldap://remote_server:389/dc=example,dc=com
```

2. Set the back-end state:

```
# dsconf <instance_name> backend suffix set --state referral
```

As a result, with the workaround, you can configure a referral for a suffix.

[Bugzilla:2063140](#)

The dsconf utility has no option to create fix-up tasks for the entryUUID plug-in

The **dsconf** utility does not provide an option to create fix-up tasks for the **entryUUID** plug-in. As a result, administrators cannot not use **dsconf** to create a task to automatically add **entryUUID** attributes to existing entries. As a workaround, create a task manually:

```
# ldapadd -D "cn=Directory Manager" -W -H ldap://server.example.com -x

dn: cn=entryuuid_fixup_<time_stamp>,cn=entryuuid task,cn=tasks,cn=config
```



```
objectClass: top
objectClass: extensibleObject
basedn: <fixup base tree>
cn: entryuuid_fixup_<time_stamp>
filter: <filtered_entry>
```

After the task has been created, Directory Server fixes entries with missing or invalid **entryUUID** attributes.

[Bugzilla:2047175](#)

MIT Kerberos does not support ECC certificates for PKINIT

MIT Kerberos does not implement the RFC5349 request for comments document, which describes the design of elliptic-curve cryptography (ECC) support in Public Key Cryptography for initial authentication (PKINIT). Consequently, the MIT **krb5-pkinit** package, used by RHEL, does not support ECC certificates. For more information, see [Elliptic Curve Cryptography \(ECC\) Support for Public Key Cryptography for Initial Authentication in Kerberos \(PKINIT\)](#).

[Bugzilla:2106043](#)

The DEFAULT:SHA1 sub-policy has to be set on RHEL 9 clients for PKINIT to work against AD KDCs

The SHA-1 digest algorithm has been deprecated in RHEL 9, and CMS messages for Public Key Cryptography for initial authentication (PKINIT) are now signed with the stronger SHA-256 algorithm.

However, the Active Directory (AD) Kerberos Distribution Center (KDC) still uses the SHA-1 digest algorithm to sign CMS messages. As a result, RHEL 9 Kerberos clients fail to authenticate users by using PKINIT against an AD KDC.

To work around the problem, enable support for the SHA-1 algorithm on your RHEL 9 systems with the following command:

```
# update-crypto-policies --set DEFAULT:SHA1
```

[Bugzilla:2060798](#)

The PKINIT authentication of a user fails if a RHEL 9 Kerberos agent communicates with a non-RHEL-9 and non-AD Kerberos agent

If a RHEL 9 Kerberos agent, either a client or Kerberos Distribution Center (KDC), interacts with a non-RHEL-9 Kerberos agent that is not an Active Directory (AD) agent, the PKINIT authentication of the user fails. To work around the problem, perform one of the following actions:

- Set the RHEL 9 agent's crypto-policy to **DEFAULT:SHA1** to allow the verification of SHA-1 signatures:

```
# update-crypto-policies --set DEFAULT:SHA1
```

- Update the non-RHEL-9 and non-AD agent to ensure it does not sign CMS data using the SHA-1 algorithm. For this, update your Kerberos client or KDC packages to the versions that use SHA-256 instead of SHA-1:
 - CentOS 9 Stream: krb5-1.19.1-15
 - RHEL 8.7: krb5-1.18.2-17

- RHEL 7.9: krb5-1.15.1-53
- Fedora Rawhide/36: krb5-1.19.2-7
- Fedora 35/34: krb5-1.19.2-3

As a result, the PKINIT authentication of the user works correctly.

Note that for other operating systems, it is the krb5-1.20 release that ensures that the agent signs CMS data with SHA-256 instead of SHA-1.

See also [The DEFAULT:SHA1 sub-policy has to be set on RHEL 9 clients for PKINIT to work against AD KDCs](#).

[Bugzilla:2077450](#)

FIPS support for AD trust requires the AD-SUPPORT crypto sub-policy

Active Directory (AD) uses AES SHA-1 HMAC encryption types, which are not allowed in FIPS mode on RHEL 9 by default. If you want to use RHEL 9 IdM hosts with an AD trust, enable support for AES SHA-1 HMAC encryption types before installing IdM software.

Since FIPS compliance is a process that involves both technical and organizational agreements, consult your FIPS auditor before enabling the **AD-SUPPORT** sub-policy to allow technical measures to support AES SHA-1 HMAC encryption types, and then install RHEL IdM:

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

[Bugzilla:2057471](#)

Heimdal client fails to authenticate a user using PKINIT against RHEL 9 KDC

By default, a Heimdal Kerberos client initiates the PKINIT authentication of an IdM user by using Modular Exponential (MODP) Diffie-Hellman Group 2 for Internet Key Exchange (IKE). However, the MIT Kerberos Distribution Center (KDC) on RHEL 9 only supports MODP Group 14 and 16.

Consequently, the pre-authentication request fails with the **krb5_get_init_creds: PREAUTH_FAILED** error on the Heimdal client and **Key parameters not accepted** on the RHEL MIT KDC.

To work around this problem, ensure that the Heimdal client uses MODP Group 14. Set the **pkinit_dh_min_bits** parameter in the **libdefaults** section of the client configuration file to 1759:

```
[libdefaults]
pkinit_dh_min_bits = 1759
```

As a result, the Heimdal client completes the PKINIT pre-authentication against the RHEL MIT KDC.

[Bugzilla:2106296](#)

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:2124243](#)

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:2060421](#)

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:2089907](#)

Users without SIDs cannot log in to IdM after an upgrade

After upgrading your IdM replica to RHEL 9.2, the IdM Kerberos Distribution Centre (KDC) might fail to issue ticket-granting tickets (TGTs) to users who do not have Security Identifiers (SIDs) assigned to their accounts. Consequently, the users cannot log in to their accounts.

To work around the problem, generate SIDs by running the following command as an IdM administrator on another IdM replica in the topology:

```
# ipa config-mod --enable-sid --add-sids
```

Afterward, if users still cannot log in, examine the Directory Server error log. You might have to adjust ID ranges to include user POSIX identities.

See the [When upgrading to RHEL9, IDM users are not able to login anymore](#) Knowledgebase solution for more information.

Jira:RHELPLAN-157939

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

MIT krb5 user fails to obtain an AD TGT because of incompatible encryption types generating the user PAC

In MIT **krb5 1.20** and later packages, a Privilege Attribute Certificate (PAC) is included in all Kerberos tickets by default. The MIT Kerberos Distribution Center (KDC) selects the strongest encryption type available to generate the KDC checksum in the PAC, which currently is the **AES HMAC-SHA2**

encryption types defined in RFC8009. However, Active Directory (AD) does not support this RFC. Consequently, in an AD-MIT cross-realm setup, an MIT **krb5** user fails to obtain an AD ticket-granting ticket (TGT) because the cross-realm TGT generated by MIT KDC contains an incompatible KDC checksum type in the PAC.

To work around the problem, set the **disable_pac** parameter to **true** for the MIT realm in the **[realms]** section of the **/var/kerberos/krb5kdc/kdc.conf** configuration file. As a result, the MIT KDC generates tickets without PAC, which means that AD skips the failing checksum verification and an MIT **krb5** user can obtain an AD TGT.

[Bugzilla:2016312](#)

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

Adding a RHEL 9 replica in FIPS mode to an IdM deployment in FIPS mode that was initialized with RHEL 8.6 or earlier fails

The default RHEL 9 FIPS cryptographic policy aiming to comply with FIPS 140-3 does not allow the use of the AES HMAC-SHA1 encryption types' key derivation function as defined by RFC3961, section 5.1.

This constraint is a blocker when adding a RHEL 9 Identity Management (IdM) replica in FIPS mode to a RHEL 8 IdM environment in FIPS mode in which the first server was installed on a RHEL 8.6 system or earlier. This is because there are no common encryption types between RHEL 9 and the previous RHEL versions, which commonly use the AES HMAC-SHA1 encryption types but do not use the AES HMAC-SHA2 encryption types.

You can view the encryption type of your IdM master key by entering the following command on the server:

```
# kadmin.local getprinc K/M | grep -E '^Key:'
```

To work around the problem, enable the use of AES HMAC-SHA1 on the RHEL 9 replica:

```
update-crypto-policies --set FIPS:AD-SUPPORT
```

WARNING

This workaround might violate FIPS compliance.

As a result, adding the RHEL 9 replica to the IdM deployment proceeds correctly.

Note that there is ongoing work to provide a procedure to generate missing AES HMAC-SHA2-

encrypted Kerberos keys on RHEL 7 and RHEL 8 servers. This will achieve FIPS 140-3 compliance on the RHEL 9 replica. However, this process will not be fully automated, because the design of Kerberos key cryptography makes it impossible to convert existing keys to different encryption types. The only way is to ask users to renew their passwords.

[Bugzilla:2103327](#)

SSSD registers the DNS names properly

Previously, if the DNS was set up incorrectly, SSSD always failed the first attempt to register the DNS name. To work around the problem, this update provides a new parameter **dns_resolver_use_search_list**. Set **dns_resolver_use_search_list = false** to avoid using the DNS search list.

[Bugzilla:1608496](#)

Directory Server terminates unexpectedly when started in referral mode

Due to a bug, global referral mode does not work in Directory Server. If you start the **ns-slaped** process with the **refer** option as the **dirsrv** user, Directory Server ignores the port settings and terminates unexpectedly. Trying to run the process as the **root** user changes SELinux labels and prevents the service from starting in future in normal mode. There are no workarounds available.

[Bugzilla:2053204](#)

Directory Server can import LDIF files only from `/var/lib/dirsrv/slaped-instance_name/ldif/`

Since RHEL 8.3, Red Hat Directory Server (RHDS) uses its own private directories and the **PrivateTmp** systemd directive is enabled by default for the LDAP services. As a result, RHDS can only import LDIF files from the `/var/lib/dirsrv/slaped-instance_name/ldif/` directory. If the LDIF file is stored in a different directory, such as `/var/tmp`, `/tmp`, or `/root`, the import fails with an error similar to the following:

```
Could not open LDIF file "/tmp/example.ldif", errno 2 (No such file or directory)
```

To work around this problem, complete the following steps:

1. Move the LDIF file to the `/var/lib/dirsrv/slaped-instance_name/ldif/` directory:

```
# mv /tmp/example.ldif /var/lib/dirsrv/slaped-instance_name/ldif/
```

2. Set permissions that allow the **dirsrv** user to read the file:

```
# chown dirsrv /var/lib/dirsrv/slaped-instance_name/ldif/example.ldif
```

3. Restore the SELinux context:

```
# restorecon -Rv /var/lib/dirsrv/slaped-instance_name/ldif/
```

For more information, see the solution article [LDAP Service cannot access files under the host's /tmp and /var/tmp directories](#).

[Bugzilla:2075525](#)

Installing a RHEL 7 IdM client with a RHEL 9.2+ IdM server in FIPS mode fails due to EMS enforcement

The TLS **Extended Master Secret** (EMS) extension (RFC 7627) is now mandatory for TLS 1.2 connections on FIPS-enabled RHEL 9.2 and later systems. This is in accordance with FIPS-140-3 requirements. However, the **openssl** version available in RHEL 7.9 and lower does not support EMS. In consequence, installing a RHEL 7 Identity Management (IdM) client with a FIPS-enabled IdM server running on RHEL 9.2 and later fails.

If upgrading the host to RHEL 8 before installing an IdM client on it is not an option, work around the problem by removing the requirement for EMS usage on the RHEL 9 server by applying a NO-ENFORCE-EMS subpolicy on top of the FIPS crypto policy:

```
# update-crypto-policies --set FIPS:NO-ENFORCE-EMS
```

Note that this removal goes against the FIPS 140-3 requirements. As a result, you can establish and accept TLS 1.2 connections that do not use EMS, and the installation of a RHEL 7 IdM client succeeds.

[Bugzilla:2220915](#)

11.13. DESKTOP

Firefox add-ons are disabled after upgrading to RHEL 9

If you upgrade from RHEL 8 to RHEL 9, all add-ons that you previously enabled in Firefox are disabled.

To work around the problem, manually reinstall or update the add-ons. As a result, the add-ons are enabled as expected.

[Bugzilla:2013247](#)

VNC is not running after upgrading to RHEL 9

After upgrading from RHEL 8 to RHEL 9, the VNC server fails to start, even if it was previously enabled.

To work around the problem, manually enable the **vncserver** service after the system upgrade:

```
# systemctl enable --now vncserver@:port-number
```

As a result, VNC is now enabled and starts after every system boot as expected.

[Bugzilla:2060308](#)

User Creation screen is unresponsive

When installing RHEL using a graphical user interface, the User Creation screen is unresponsive. As a consequence, creating users during installation is more difficult.

To work around this problem, use one of the following solutions to create users:

- Run the installation in VNC mode and resize the VNC window.
- Create users after completing the installation process.

[BZ#2122636](#)

11.14. GRAPHICS INFRASTRUCTURES

NVIDIA drivers might revert to X.org

Under certain conditions, the proprietary NVIDIA drivers disable the Wayland display protocol and revert to the X.org display server:

- If the version of the NVIDIA driver is lower than 470.
- If the system is a laptop that uses hybrid graphics.
- If you have not enabled the required NVIDIA driver options.

Additionally, Wayland is enabled but the desktop session uses X.org by default if the version of the NVIDIA driver is lower than 510.

Jira:RHELPLAN-119001

Night Light is not available on Wayland with NVIDIA

When the proprietary NVIDIA drivers are enabled on your system, the **Night Light** feature of GNOME is not available in Wayland sessions. The NVIDIA drivers do not currently support **Night Light**.

Jira:RHELPLAN-119852

Matrox G200e shows no output on a VGA display

Your display might show no graphical output if you use the following system configuration:

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

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

To work around the problem, use the following procedure:

1. Boot the system to the boot loader menu.
2. Add the **module_blacklist=mgag200** option to the kernel command line.

As a result, RHEL boots and shows graphical output as expected, but the maximum resolution is limited to 1024x768 at the 16-bit color depth.

Bugzilla:1960467

X.org configuration utilities do not work under Wayland

X.org utilities for manipulating the screen do not work in the Wayland session. Notably, the **xrandr** utility does not work under Wayland due to its different approach to handling, resolutions, rotations, and layout.

Jira:RHELPLAN-121049

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

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:2078999](#)

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](#)

Unable to register systems with environment names

The **rhc** system role fails to register the system when specifying environment names in **rhc_environment**. As a workaround, use environment IDs instead of environment names while registering.

[Bugzilla:2187539](#)

11.17. VIRTUALIZATION

Installing a virtual machine over https or ssh in some cases fails

Currently, the **virt-install** utility fails when attempting to install a guest operating system (OS) from an ISO source over a https or ssh connection - for example using **virt-install --cdrom https://example/path/to/image.iso**. Instead of creating a virtual machine (VM), the described operation terminates unexpectedly with an **internal error: process exited while connecting to monitor** message.

Similarly, using the RHEL 9 web console to install a guest OS fails and displays an **Unknown driver 'https'** error if you use an https or ssh URL, or the **Download OS** function.

To work around this problem, install **qemu-kvm-block-curl** and **qemu-kvm-block-ssh** on the host to enable https and ssh protocol support, respectively. Alternatively, use a different connection protocol or a different installation source.

[Bugzilla:2014229](#)

Using NVIDIA drivers in virtual machines disables Wayland

Currently, NVIDIA drivers are not compatible with the Wayland graphical session. As a consequence, RHEL guest operating systems that use NVIDIA drivers automatically disable Wayland and load an Xorg session instead. This primarily occurs in the following scenarios:

- When you pass through an NVIDIA GPU device to a RHEL virtual machine (VM)
- When you assign an NVIDIA vGPU mediated device to a RHEL VM

Jira:RHELPLAN-117234

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 (**fsrm**) 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 **fsrm** in the BIOS of your host.

Bugzilla:2077767

A hostdev interface with failover settings cannot be hot-plugged after being hot-unplugged

After removing a **hostdev** network interface with failover configuration from a running virtual machine (VM), the interface currently cannot be re-attached to the same running VM.

Bugzilla:2052424

Live post-copy migration of VMs with failover VFs fails

Currently, attempting to post-copy migrate a running virtual machine (VM) fails if the VM uses a device with the virtual function (VF) failover capability enabled. To work around the problem, use the standard migration type, rather than post-copy migration.

Bugzilla:1817965

Host network cannot ping VMs with VFs during live migration

When live migrating a virtual machine (VM) with a configured virtual function (VF), such as a VMs that uses virtual SR-IOV software, the network of the VM is not visible to other devices and the VM cannot be reached by commands such as **ping**. After the migration is finished, however, the problem no longer occurs.

Bugzilla:1789206

Failover virtio NICs are not assigned an IP address on Windows virtual machines

Currently, when starting a Windows virtual machine (VM) with only a failover virtio NIC, the VM fails to assign an IP address to the NIC. Consequently, the NIC is unable to set up a network connection. Currently, there is no workaround.

Bugzilla:1969724

Disabling AVX causes VMs to become unbootable

On a host machine that uses a CPU with Advanced Vector Extensions (AVX) support, attempting to boot a VM with AVX explicitly disabled currently fails, and instead triggers a kernel panic in the VM.

Bugzilla:2005173

Windows VM fails to get IP address after network interface reset

Sometimes, Windows virtual machines fail to get an IP address after an automatic network interface reset. As a consequence, the VM fails to connect to the network. To work around this problem, disable and re-enable the network adapter driver in the Windows Device Manager.

Bugzilla:2084003

Broadcom network adapters work incorrectly on Windows VMs after a live migration

Currently, network adapters from the Broadcom family of devices, such as Broadcom, Qlogic, or Marvell, cannot be hot-unplugged during live migration of Windows virtual machines (VMs). As a consequence, the adapters work incorrectly after the migration is complete.

This problem affects only those adapters that are attached to Windows VMs using Single-root I/O virtualization (SR-IOV).

[Bugzilla:2090712](#), [Bugzilla:2091528](#), [Bugzilla:2111319](#)

Windows Server 2016 VMs sometimes stops working after hot-plugging a vCPU

Currently, assigning a vCPU to a running virtual machine (VM) with a Windows Server 2016 guest operating system might cause a variety of problems, such as the VM terminating unexpectedly, becoming unresponsive, or rebooting.

[Bugzilla:1915715](#)

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:2020146](#)

Redundant error messages on VMs with NVIDIA passthrough devices

When using an Intel host machine with a RHEL 9.2 operating system, virtual machines (VMs) with a passed through NVIDIA GPU device frequently log the following error message:

Spurious APIC interrupt (vector 0xFF) on CPU#2, should never happen.

However, this error message does not impact the functionality of the VM and can be ignored. For details, see the [Red Hat KnowledgeBase](#).

[Bugzilla:2149989](#)

Some Windows guests fail to boot after a v2v conversion on hosts with AMD EPYC CPUs

After using the **virt-v2v** utility to convert a virtual machine (VM) that uses Windows 11 or a Windows Server 2022 as the guest OS, the VM currently fails to boot. This occurs on hosts that use AMD EPYC series CPUs.

[Bugzilla:2168082](#)

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:1947422](#)

The Nvidia GPU driver stops working after the VM shutdown

The RHEL kernel has adopted an upstream Linux change that aligns device power transitions delays more closely to those required by the PCIe specification. As a consequence, due to the audio function of the GPU, some Nvidia GPUs might stop working after the shutdown of a VM.

To work around the problem, unassign the audio function of the GPU from the VM. In addition, due to the DMA isolation requirements for device assignment (that is, IOMMU grouping), bind the audio function to the **vfio-pci** driver, which allows the GPU function to continue to be assigned and function normally.

[Bugzilla:2178956](#)

nodedev-dumpxml does not list attributes correctly for certain mediated devices

Currently, the **nodedev-dumpxml** does not list attributes correctly for mediated devices that were created using the **nodedev-create** command. To work around this problem, use the **nodedev-define** and **nodedev-start** commands instead.

[Bugzilla:2143158](#)

Recovering an interrupted post-copy VM migration might fail

If a post-copy migration of a virtual machine (VM) is interrupted and then immediately resumed on the same incoming port, the migration might fail with the following error: **Address already in use**

To work around this problem, wait at least 10 seconds before resuming the post-copy migration or switch to another port for migration recovery.

[Bugzilla:2178376](#)

virtiofs devices cannot be attached after restarting virtqemud or libvirtd

Currently, restarting the **virtqemud** or **libvirtd** services prevents **virtiofs** storage devices from being attached to virtual machines on your host.

[Bugzilla:2078693](#)

virsh blkio tune --weight command fails to set the correct cgroup I/O controller value

Currently, using the **virsh blkio tune --weight** command to set the VM weight does not work as expected. The command fails to set the correct **io.bfq.weight** value in the cgroup I/O controller interface file. There is no workaround at this time.

[Jira:RHELPLAN-83423](#)

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](#)

NUMA node mapping not working correctly on AMD EPYC CPUs

QEMU does not handle NUMA node mapping on AMD EPYC CPUs correctly. As a result, the performance of virtual machines (VMs) with these CPUs might be negatively impacted if using a NUMA node configuration. In addition, the VMs display a warning similar to the following during boot.

```
sched: CPU #4's llc-sibling CPU #3 is not on the same node! [node: 1 != 0]. Ignoring dependency.
WARNING: CPU: 4 PID: 0 at arch/x86/kernel/smpboot.c:415 topology_sane.isra.0+0x6b/0x80
```

To work around this issue, do not use AMD EPYC CPUs for NUMA node configurations.

[Bugzilla:2176010](#)

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:2058982](#)

PCIe ATS devices do not work on Windows VMs

When you configure a PCIe Address Translation Services (ATS) device in the XML configuration of virtual machine (VM) with a Windows guest operating system, the guest does not enable the ATS device after booting the VM. This is because Windows currently does not support ATS on **virtio** devices.

[Bugzilla:2073872](#)

11.18. RHEL IN CLOUD ENVIRONMENTS

Cloning or restoring RHEL 9 virtual machines that use LVM on Nutanix AHV causes non-root partitions to disappear

When running a RHEL 9 guest operating system on a virtual machine (VM) hosted on the Nutanix AHV hypervisor, restoring the VM from a snapshot or cloning the VM currently causes non-root partitions in the VM to disappear if the guest is using Logical Volume Management (LVM). As a consequence, the following problems occur:

- After restoring the VM from a snapshot, the VM cannot boot, and instead enters emergency mode.
- A VM created by cloning cannot boot, and instead enters emergency mode.

To work around these problems, do the following in emergency mode of the VM:

1. Remove the LVM system devices file: **rm /etc/lvm/devices/system.devices**

2. Recreate LVM device settings: **vgimportdevices -a**
3. Reboot the VM

This makes it possible for the cloned or restored VM to boot up correctly.

Alternatively, to prevent the issue from occurring, do the following before cloning a VM or creating a VM snapshot:

1. Uncomment the **use_devicesfile = 0** line in the `/etc/lvm/lvm.conf` file
2. Reboot the VM

Bugzilla:2059545

Customizing RHEL 9 guests on ESXi sometimes causes networking problems

Currently, customizing a RHEL 9 guest operating system in the VMware ESXi hypervisor does not work correctly with NetworkManager key files. As a consequence, if the guest is using such a key file, it will have incorrect network settings, such as the IP address or the gateway.

For details and workaround instructions, see the [VMware Knowledge Base](#).

Bugzilla:2037657

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

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](#)

11.19. SUPPORTABILITY

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:1869561

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

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:2096795 , Bugzilla:1859271 , Bugzilla:2057070 , Bugzilla:2093981 , Bugzilla:1132524 , Bugzilla:2136610 , Bugzilla:2142639 , Bugzilla:1878808 , Bugzilla:1924569 , Bugzilla:1956987 , Bugzilla:1952241 , Bugzilla:2063140 , Bugzilla:2047175 , Bugzilla:2053204
Doc-administration-guide	Bugzilla:2075525
NetworkManager	Bugzilla:2134897 , Bugzilla:2081302 , Bugzilla:2019306 , Bugzilla:2128809 , Bugzilla:2110307 , Bugzilla:2117352 , Bugzilla:2029636 , Bugzilla:2073512 , Bugzilla:2128216 , Bugzilla:1894877 , Bugzilla:2151040
aardvark-dns	Jira:RHELPLAN-138024
anaconda	Bugzilla:2052938 , Bugzilla:2158210 , Bugzilla:1991843 , Bugzilla:2127100 , Bugzilla:2093793 , Bugzilla:2107346 , Bugzilla:2050140 , Bugzilla:1877697 , Bugzilla:1914955 , Bugzilla:1929105 , Bugzilla:1997832 , Bugzilla:2125542 , Bugzilla:2115783 , Bugzilla:2164216 , Bugzilla:2163497
ansible-collection-microsoft-sql	Bugzilla:2151282 , Bugzilla:2151283 , Bugzilla:2151284 , Bugzilla:2153428 , Bugzilla:2163709
ansible-freeipa	Bugzilla:2127913
bacula	Bugzilla:2089395
bind	Bugzilla:1984982
chrony	Bugzilla:2133754
clevis	Bugzilla:2126533 , Bugzilla:2159728 , Bugzilla:2159735
cloud-init	Bugzilla:1750862
cockpit	Bugzilla:2207498
cockpit-appstream	Bugzilla:2030836
cockpit-machines	Bugzilla:2173584
conntrack-tools	Bugzilla:2132398

Component	Tickets
crash	Bugzilla:2119685
crypto-policies	Bugzilla:2152635
cyrus-sasl	Bugzilla:1995600
device-mapper-multipath	Bugzilla:2033080 , Bugzilla:2011699 , Bugzilla:1926147
dnf	Bugzilla:2131288 , Bugzilla:2121662 , Bugzilla:2122626 , Bugzilla:2073510
dnf-plugins-core	Bugzilla:2139326
edk2	Bugzilla:1935497
fapolicyd	Jira:RHEL-192 , Bugzilla:2054740 , Bugzilla:2070655
firefox	Bugzilla:2013247
firewalld	Bugzilla:2125371 , Bugzilla:2077512 , Bugzilla:2122678
frr	Bugzilla:2129731 , Bugzilla:2129743
gcc	Bugzilla:2110583 , Bugzilla:2117632 , Bugzilla:2141718
gdm	Bugzilla:2131203
gimp	Bugzilla:2047161
git	Bugzilla:2139379
git-lfs	Bugzilla:2139383
glibc	Bugzilla:2129005 , Bugzilla:2155352
gnome-shell-extensions	Bugzilla:2154358 , Bugzilla:2160553
gnupg2	Bugzilla:2070722 , Bugzilla:2073567
gnutls	Bugzilla:2084161 , Bugzilla:2042009
golang	Bugzilla:2133019 , Bugzilla:2175173 , Bugzilla:2111072 , Bugzilla:2092016
grafana	Bugzilla:2116847
grafana-pcp	Bugzilla:2116848

Component	Tickets
grub2	Bugzilla:2026579
grubby	Bugzilla:2127453
gssproxy	Bugzilla:2184333
ipa	Bugzilla:2143224 , Bugzilla:2162677 , Bugzilla:2084180 , Bugzilla:2084166 , Bugzilla:2069202 , Bugzilla:2094673 , Bugzilla:2057471 , Bugzilla:2124243 , Bugzilla:2089907
iproute	Bugzilla:2155604
java-1.8.0-openjdk	Bugzilla:2188023
java-17-openjdk	Bugzilla:2186803 , Bugzilla:2186810 , Bugzilla:2186806
jmc	Bugzilla:2122401
jmc-core	Bugzilla:1980981
kdump-anaconda-addon	Bugzilla:2017401
kernel	Bugzilla:2153073 , Bugzilla:2143850 , Bugzilla:1871126 , Bugzilla:1871143 , Bugzilla:2075216 , Bugzilla:2100606 , Bugzilla:2104468 , Bugzilla:2111048 , Bugzilla:2150284 , Bugzilla:2066372 , Bugzilla:2107347 , Bugzilla:2140899 , Bugzilla:2069758 , Bugzilla:1613522 , Bugzilla:1874182 , Bugzilla:1995338 , Bugzilla:1570255 , Bugzilla:2023416 , Bugzilla:2021672 , Bugzilla:2027304 , Bugzilla:1660337 , Bugzilla:1955275 , Bugzilla:2142102 , Bugzilla:2041690 , Bugzilla:2040643 , Bugzilla:2167783 , Bugzilla:2000616 , Bugzilla:2013650 , Bugzilla:2132480 , Bugzilla:2059545 , Bugzilla:1960467 , Bugzilla:2005173 , Bugzilla:2128610 , Bugzilla:2129288 , Bugzilla:2013884 , Bugzilla:2149989 , Bugzilla:2168603 , Bugzilla:2173947 , Bugzilla:2178956 , Bugzilla:2180665
kexec-tools	Bugzilla:2085347 , Bugzilla:2076416 , Bugzilla:2160676 , Bugzilla:2080110 , Bugzilla:2139000 , Bugzilla:2113873 , Bugzilla:2064708 , Bugzilla:2065013
keylime	Bugzilla:2150830 , Bugzilla:2138167 , Bugzilla:2140670 , Bugzilla:2142009
kmod	Bugzilla:2103605
krb5	Bugzilla:2068535 , Bugzilla:2106043 , Bugzilla:2060798 , Bugzilla:2077450 , Bugzilla:2106296 , Bugzilla:2060421 , Bugzilla:2016312 , Bugzilla:2103327
libdnf	Bugzilla:2124480
libnvme	Bugzilla:2139752

Component	Tickets
libotr	Bugzilla:2086562
libreswan	Bugzilla:2128669
libsepol	Bugzilla:2145224
libssh	Bugzilla:2026449 , Bugzilla:2068475
libvirt	Bugzilla:2014487 , Bugzilla:2143158 , Bugzilla:2078693
libxcrypt	Bugzilla:2034569
llvm-toolset	Bugzilla:2118567
lvm2	Bugzilla:1878893 , Bugzilla:2038183
mod_security	Bugzilla:2143211
mysql	Bugzilla:1991500
nfs-utils	Bugzilla:2143747 , Bugzilla:2081114
nginx	Bugzilla:2096174
nmstate	Bugzilla:2095207 , Bugzilla:2120473 , Bugzilla:2044150 , Bugzilla:2058292 , Bugzilla:2130240 , Bugzilla:2162401
nodejs	Bugzilla:2178088
nss	Bugzilla:2091905
nvme-cli	Bugzilla:2139753
nvme-stas	Bugzilla:1893841
open-vm-tools	Bugzilla:2037657
openblas	Bugzilla:2112099 , Bugzilla:2115737
opencryptoki	Bugzilla:2110314
openscap	Bugzilla:2159286 , Bugzilla:2161499
openslp	Bugzilla:2184570

Component	Tickets
openssh	Bugzilla:2056884
openssl	Bugzilla:2129063 , Bugzilla:2188046 , Bugzilla:2060044 , Bugzilla:1975836 , Bugzilla:2168665 , Bugzilla:1681178 , Bugzilla:1685470
openssl-ibmca	Bugzilla:2110378
osbuild-composer	Bugzilla:2173928
oscap-anaconda-addon	Bugzilla:2165920 , Bugzilla:2172264
pacemaker	Bugzilla:2133546 , Bugzilla:2125344 , Bugzilla:2125337
pam	Bugzilla:2126640
passt	Bugzilla:2131015
pause-container	Bugzilla:2106816
pcp	Bugzilla:2117074
pcs	Bugzilla:2116295 , Bugzilla:2112270 , Bugzilla:1620043 , Bugzilla:1796827 , Bugzilla:2092950
pki-core	Bugzilla:1849834 , Bugzilla:1883477
podman	Jira:RHELPLAN-136602 , Jira:RHELPLAN-136607 , Bugzilla:2119200 , Jira:RHELPLAN-136611 , Bugzilla:2069279
postgresql	Bugzilla:2128410
powerpc-utils	Bugzilla:2125152
powertop	Bugzilla:2044132
python-blivet	Bugzilla:2103800
python-sqlalchemy	Bugzilla:2152649
python3.11	Bugzilla:2127923
python3.11-lxml	Bugzilla:2157708

Component	Tickets
qemu-kvm	Bugzilla:2116496 , Bugzilla:1965079 , Bugzilla:1951814 , Bugzilla:2060839 , Bugzilla:2014229 , Bugzilla:2052424 , Bugzilla:1817965 , Bugzilla:1789206 , Bugzilla:2090712 , Bugzilla:1915715 , Bugzilla:2020146 , Bugzilla:1947422 , Bugzilla:2178376 , Bugzilla:2176010 , Bugzilla:2058982
realtime-tests	Bugzilla:2041637
rear	Bugzilla:2172589 , Bugzilla:2160748
restore	Bugzilla:1997366
rhel-system-roles	Bugzilla:2131293 , Bugzilla:2133858 , Bugzilla:2078999 , Bugzilla:2119102 , Bugzilla:2128843 , Bugzilla:2130010 , Bugzilla:2130329 , Bugzilla:2130344 , Bugzilla:2130357 , Bugzilla:2133528 , Bugzilla:2133930 , Bugzilla:2134202 , Bugzilla:2137663 , Bugzilla:2140795 , Bugzilla:2141330 , Bugzilla:2143768 , Bugzilla:2165175 , Bugzilla:2140804 , Bugzilla:2126959 , Bugzilla:2143816 , Bugzilla:2153030 , Bugzilla:2153043 , Bugzilla:2162782 , Bugzilla:2167528 , Bugzilla:2168735 , Bugzilla:2160152 , Bugzilla:1999770 , Bugzilla:2123859 , Bugzilla:2187539 , Bugzilla:2186218
rpm	Bugzilla:2150804 , Bugzilla:2111251 , Bugzilla:2144005
rsyslog	Bugzilla:2124849 , Bugzilla:2127404 , Bugzilla:2124488 , Bugzilla:2157659
rteval	Bugzilla:2081325
rust	Bugzilla:2123900
s390utils	Bugzilla:2044204 , Bugzilla:1932480
samba	Bugzilla:2131993 , Jira:RHELDPCS-16612
scap-security-guide	Bugzilla:2158405 , Bugzilla:2122325 , Bugzilla:2169414 , Bugzilla:2105162 , Bugzilla:2120978 , Bugzilla:2038978
selinux-policy	Bugzilla:2151841 , Bugzilla:1972222 , Bugzilla:2064274
sos	Bugzilla:2164987 , Bugzilla:2134906 , Bugzilla:1869561
sssd	Bugzilla:1507035 , Bugzilla:2087247 , Bugzilla:1766490 , Bugzilla:2065693 , Bugzilla:2056482 , Bugzilla:1608496
stratisd	Bugzilla:2039957 , Bugzilla:2039955 , Bugzilla:2041558
subscription-manager	Bugzilla:2108549 , Bugzilla:2163716 , Bugzilla:2136694

Component	Tickets
swig	Bugzilla:2139101
synce4l	Bugzilla:2143264
systemd	Bugzilla:2217931 , Bugzilla:2018112
systemtap	Bugzilla:2083727
tang	Bugzilla:2095474 , Bugzilla:2188743
tigervnc	Bugzilla:2060308
tomcat	Bugzilla:2160511
toolbox	Bugzilla:2163752
tuna	Bugzilla:2122781 , Bugzilla:2121517 , Bugzilla:2062865
tuned	Bugzilla:2133815 , Bugzilla:2113900
tzdata	Bugzilla:2157982
udisks2	Bugzilla:1983602
unbound	Bugzilla:2070495
usbguard	Bugzilla:2155910 , Bugzilla:2042345 , Bugzilla:2097419
virt-v2v	Bugzilla:2168082
virtio-win	Bugzilla:1969724 , Bugzilla:2084003
vsftpd	Bugzilla:2018284
wsmancli	Bugzilla:2127416
xdp-tools	Bugzilla:2160066

Component	Tickets
other	<p> Bugzilla:2177782, Jira:RHELPLAN-137505, Jira:RHELPLAN-139125, Bugzilla:2046653, Jira:RHELPLAN-133650, Jira:RHELPLAN-139430, Jira:RHELPLAN-137416, Jira:RHELPLAN-137411, Jira:RHELPLAN-137406, Jira:RHELPLAN-137403, Jira:RHELPLAN-159146, Jira:RHELPLAN-139448, Jira:RHELPLAN-151481, Jira:RHELPLAN-150266, Jira:RHELPLAN-147982, Jira:RHELPLAN-147428, Jira:RHELPLAN-139659, Jira:RHELPLAN-149091, Jira:RHELPLAN-139655, Jira:RHELPLAN-139424, Jira:RHELPLAN-136489, Jira:RHELPLAN-59528, Bugzilla:2209419, Bugzilla:2190123, Jira:RHELPLAN-135600, Jira:RHELPLAN-148303, Bugzilla:2020529, Bugzilla:2030412, Jira:RHELPLAN-103993, Jira:RHELPLAN-122345, Jira:RHELPLAN-27394, Jira:RHELPLAN-27737, Jira:RHELPLAN-148394, Bugzilla:1927780, Jira:RHELPLAN-110763, Bugzilla:1935544, Bugzilla:2089200, Jira:RHELPLAN-15509, Jira:RHELPLAN-99136, Jira:RHELPLAN-103232, Bugzilla:1899167, Bugzilla:1979521, Jira:RHELPLAN-100087, Jira:RHELPLAN-100639, Bugzilla:2058153, Jira:RHELPLAN-113995, Jira:RHELPLAN-98983, Jira:RHELPLAN-131882, Jira:RHELPLAN-137660, Jira:RHELPLAN-139805, Jira:RHELPLAN-147725, Jira:RHELPLAN-153267, Jira:RHELDOS-16300, Jira:RHELPLAN-157225, Jira:RHELPLAN-157337, Bugzilla:1640697, Bugzilla:1697896, Bugzilla:2047713, Jira:RHELPLAN-96940, Jira:RHELPLAN-117234, Jira:RHELPLAN-119001, Jira:RHELPLAN-119852, Bugzilla:2077767, Bugzilla:2053598, Bugzilla:2082303, Jira:RHELPLAN-121049, Jira:RHELPLAN-157939, Jira:RHELPLAN-109613, Bugzilla:2160619, Bugzilla:2173992, Bugzilla:2185048, Jira:RHELPLAN-83423 </p>

APPENDIX B. REVISION HISTORY

0.2-1

September 25 2023, Gabriela Fialová (gfialova@redhat.com)

- Added a known issue [BZ#2122636](#) (Desktop).

0.2-0

September 13 2023, Lenka Špačková (lspackova@redhat.com)

- Fixed command formatting in [BZ#2220915](#).

0.1-9

September 8 2023, Marc Muehlfeld (mmuehlfeld@redhat.com)

- Added a deprecated functionality release note [JIRA:RHELDPCS-16612](#) (Samba).
- Updated the [Providing feedback on Red Hat documentation](#) section.

0.1-8

September 5 2023, Gabriela Fialová (gfialova@redhat.com)

- Added an enhancement [BZ#2075017](#) (idm_ds).

0.1-7

August 31 2023, Gabriela Fialová (gfialova@redhat.com)

- Added a known issue [BZ#2230431](#) (plumbers).

0.1-6

August 29 2023, Gabriela Fialová (gfialova@redhat.com)

- Added a known issue [BZ#2220915](#) (IdM).

0.1-5

August 25 2023, Lucie Vařáková (lvarakova@redhat.com)

- Added a known issue [BZ#2214508](#) (Kernel).

0.1.4

August 17 2023, Gabriela Fialová (gfialova@redhat.com)

- Add an enhancement [BZ#2136937](#) (Plumbers).

0.1.3

August 14 2023, Lenka Špačková (lspackova@redhat.com)

- Fixed a typo in [BZ#2128410](#).

0.1.2

August 09 2023, Gabriela Fialová (gfialova@redhat.com)

- Added a known issue [BZ#2155010](#) (COP)

- Updated a Security Bug Fix [BZ#2155910](#) (CS).

0.1.1

August 07 2023, Gabriela Fialová (gfialova@redhat.com)

- Updated a deprecated functionality release note [BZ#2214130](#) (CS).

0.1.0

August 03 2023, Lenka Špačková (lspackova@redhat.com)

- Fixed formatting in [BZ#2142639](#) and [BZ#2119102](#).
- Improved abstract.

0.0.9

August 02 2023, Marc Muehlfeld (mmuehlfeld@redhat.com)

- Updated a deprecated functionality release note [BZ#1894877](#) (NetworkManager).

0.0.8

Aug 1 2023, Mirek Jahoda (mjahoda@redhat.com)

- Replaced the web console known issue with NBDE by a bug fix [BZ#2207498](#) (RHEL web console).

0.0.7

July 27 2023, Gabriela Fialová (gfialova@redhat.com)

- Amended 3 enhancements in kernel and 1 in compilers and dev tools as per DDF feedback.

0.0.6

July 25 2023, Gabriela Fialová (gfialova@redhat.com)

- Added a Known Issue [BZ#2109231](#) (Installer).

0.0.5

June 22 2023, Gabriela Fialová (gfialova@redhat.com)

- Added an Enhancement [BZ#2087247](#) (IdM).

0.0.4

June 8 2023, Gabriela Fialová (gfialova@redhat.com)

- Added an Enhancement [BZ#2190123](#) (kernel).

0.0.3

June 6 2023, Gabriela Fialová (gfialova@redhat.com)

- Added [RHELPLAN-159146](#) (IdM).

0.0.2

June 5 2023, Gabriela Fialová (gfialova@redhat.com)

- Added a KI [BZ#2176010](#) (virt).

0.0.1

May 10 2023, Gabriela Fialová (gfialova@redhat.com)

- Release of the Red Hat Enterprise Linux 9.2 Release Notes.

0.0.0

March 29 2023, Gabriela Fialová (gfialova@redhat.com)

- Release of the Red Hat Enterprise Linux 9.2 Beta Release Notes.