



Red Hat Enterprise Linux 9.3

9.3 Release Notes

Release Notes for Red Hat Enterprise Linux 9.3

Red Hat Enterprise Linux 9.3 9.3 Release Notes

Release Notes for Red Hat Enterprise Linux 9.3

Legal Notice

Copyright © 2024 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 9.3 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

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	5
CHAPTER 1. OVERVIEW	6
1.1. MAJOR CHANGES IN RHEL 9.3	6
Installer and image creation	6
1.1.1. Bootloader	6
RHEL for Edge	6
Security	6
Dynamic programming languages, web and database servers	7
Compilers and development tools	7
Updated system toolchain	7
Updated performance tools and debuggers	7
Updated performance monitoring tools	7
Updated compiler toolsets	7
Java implementations in RHEL 9	8
1.2. IN-PLACE UPGRADE	8
In-place upgrade from RHEL 8 to RHEL 9	8
In-place upgrade from RHEL 7 to RHEL 9	9
1.3. RED HAT CUSTOMER PORTAL LABS	9
1.4. ADDITIONAL RESOURCES	10
CHAPTER 2. ARCHITECTURES	11
CHAPTER 3. DISTRIBUTION OF CONTENT IN RHEL 9	12
3.1. INSTALLATION	12
3.2. REPOSITORIES	12
3.3. APPLICATION STREAMS	13
3.4. PACKAGE MANAGEMENT WITH YUM/DNF	13
CHAPTER 4. NEW FEATURES	14
4.1. INSTALLER AND IMAGE CREATION	14
4.2. SECURITY	15
4.3. RHEL FOR EDGE	21
4.4. SOFTWARE MANAGEMENT	23
4.5. SHELLS AND COMMAND-LINE TOOLS	23
4.6. INFRASTRUCTURE SERVICES	24
4.7. NETWORKING	25
4.8. KERNEL	32
4.9. BOOT LOADER	36
4.10. FILE SYSTEMS AND STORAGE	36
4.11. HIGH AVAILABILITY AND CLUSTERS	39
4.12. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	41
4.13. COMPILERS AND DEVELOPMENT TOOLS	45
4.14. IDENTITY MANAGEMENT	63
4.15. GRAPHICS INFRASTRUCTURES	68
4.16. THE WEB CONSOLE	69
4.17. RED HAT ENTERPRISE LINUX SYSTEM ROLES	69
4.18. VIRTUALIZATION	72
4.19. RHEL IN CLOUD ENVIRONMENTS	74
4.20. SUPPORTABILITY	74
4.21. CONTAINERS	75

CHAPTER 5. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS	78
New kernel parameters	78
Updated kernel parameters	79
Removed kernel parameters	84
CHAPTER 6. DEVICE DRIVERS	85
6.1. NEW DRIVERS	85
Network drivers	85
Graphics drivers and miscellaneous drivers	87
6.2. UPDATED DRIVERS	89
Network driver updates	89
Storage driver updates	90
CHAPTER 7. AVAILABLE BPF FEATURES	91
CHAPTER 8. BUG FIXES	110
8.1. INSTALLER AND IMAGE CREATION	110
8.2. SECURITY	110
8.3. SUBSCRIPTION MANAGEMENT	116
8.4. SOFTWARE MANAGEMENT	116
8.5. SHELLS AND COMMAND-LINE TOOLS	116
8.6. NETWORKING	118
8.7. KERNEL	118
8.8. BOOT LOADER	118
8.9. FILE SYSTEMS AND STORAGE	119
8.10. HIGH AVAILABILITY AND CLUSTERS	119
8.11. COMPILERS AND DEVELOPMENT TOOLS	121
8.12. IDENTITY MANAGEMENT	122
8.13. THE WEB CONSOLE	125
8.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES	125
8.15. VIRTUALIZATION	129
CHAPTER 9. TECHNOLOGY PREVIEWS	131
9.1. INSTALLER AND IMAGE CREATION	131
9.2. SECURITY	131
9.3. SHELLS AND COMMAND-LINE TOOLS	131
9.4. INFRASTRUCTURE SERVICES	131
9.5. NETWORKING	132
9.6. KERNEL	133
9.7. FILE SYSTEMS AND STORAGE	135
9.8. COMPILERS AND DEVELOPMENT TOOLS	136
9.9. IDENTITY MANAGEMENT	137
9.10. DESKTOP	139
9.11. VIRTUALIZATION	140
9.12. RHEL IN CLOUD ENVIRONMENTS	141
9.13. CONTAINERS	141
CHAPTER 10. DEPRECATED FUNCTIONALITY	143
10.1. INSTALLER AND IMAGE CREATION	143
10.2. SECURITY	144
10.3. SUBSCRIPTION MANAGEMENT	146
10.4. SHELLS AND COMMAND-LINE TOOLS	146
10.5. NETWORKING	147
10.6. KERNEL	148

10.7. FILE SYSTEMS AND STORAGE	148
10.8. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	149
10.9. COMPILERS AND DEVELOPMENT TOOLS	150
10.10. IDENTITY MANAGEMENT	150
10.11. DESKTOP	152
10.12. GRAPHICS INFRASTRUCTURES	152
10.13. RED HAT ENTERPRISE LINUX SYSTEM ROLES	153
10.14. VIRTUALIZATION	153
10.15. CONTAINERS	155
10.16. DEPRECATED PACKAGES	155
CHAPTER 11. KNOWN ISSUES	170
11.1. INSTALLER AND IMAGE CREATION	170
11.2. SECURITY	174
11.3. RHEL FOR EDGE	178
11.4. SOFTWARE MANAGEMENT	178
11.5. SHELLS AND COMMAND-LINE TOOLS	179
11.6. INFRASTRUCTURE SERVICES	181
11.7. NETWORKING	181
11.8. KERNEL	183
11.9. FILE SYSTEMS AND STORAGE	187
11.10. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	189
11.11. IDENTITY MANAGEMENT	190
11.12. DESKTOP	194
11.13. GRAPHICS INFRASTRUCTURES	195
11.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES	195
11.15. VIRTUALIZATION	196
11.16. RHEL IN CLOUD ENVIRONMENTS	201
11.17. SUPPORTABILITY	203
11.18. CONTAINERS	203
APPENDIX A. LIST OF TICKETS BY COMPONENT	205
APPENDIX B. REVISION HISTORY	214

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

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

Submitting feedback through Jira (account required)

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

CHAPTER 1. OVERVIEW

1.1. MAJOR CHANGES IN RHEL 9.3

Installer and image creation

Key highlights for image builder:

- Enhancement to the AWS EC2 AMD or Intel 64-bit architecture AMI image to support UEFI boot, in addition to the legacy BIOS boot.

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

1.1.1. Bootloader

New default behavior of `grub2-mkconfig` with BLS

With this release, the `grub2-mkconfig` command no longer overwrites the kernel command line in Boot Loader Specification (BLS) snippets with `GRUB_CMDLINE_LINUX` by default. Each kernel in the boot loader menu takes its kernel command line from its BLS snippet. This new default behavior is caused by the `GRUB_ENABLE_BLSCFG=true` option.

For details, see [New features in Bootloader](#).

RHEL for Edge

Key highlights for RHEL for Edge:

- Support added to the following image types:
 - **minimal-raw**
 - **edge-vsphere**
 - **edge-ami**
- New FIDO Device Onboarding Servers container images available
 - `rhel9/fdo-manufacturing-server`
 - `rhel9/fdo-owner-onboarding-server`
 - `rhel9/fdo-rendezvous-server`
 - `rhel9/fdo-serviceinfo-api-server`

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

Security

Key security-related highlights:

- **Keylime** was rebased to version 7.3.0.
- The **keylime RHEL System Role** is available. With this role, you can more easily configure the Keylime verifier and Keylime registrar.
- **OpenSSH** was migrated further from the less secure SHA-1 message digest for cryptographic purposes, and instead applies the more secure SHA-2 in additional scenarios.

- The **pcsc-lite-ccid** USB Chip/Smart Card Interface Device(CCID)) and Integrated Circuit Card Device (ICCD) driver was rebased to version 1.5.2.
- RHEL 9.3 introduces further improvements to support the **Extended Master Secret (EMS)** extension (RFC 7627) required by the FIPS-140-3 standard for all TLS 1.2 connections.
- **SEtools**, the collection of graphical tools, command-line tools, and libraries for SELinux policy analysis, was rebased to version 4.4.3.
- **OpenSCAP** was rebased to version 1.3.8.
- **SCAP Security Guide** was rebased to version 0.1.69, most notably:
 - ANSSI profiles were updated to version 2.0.
 - Three new SCAP profiles were added for RHEL 9 aligned with the CCN-STIC-610A22 Guide.

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

Dynamic programming languages, web and database servers

Later versions of the following Application Streams are now available:

- **Redis 7**
- **Node.js 20**

In addition, the **Apache HTTP Server** has been updated to version 2.4.57.

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

Compilers and development tools

Updated system toolchain

The following system toolchain component has been updated in RHEL 9.3:

- **GCC 11.4.1**

Updated performance tools and debuggers

The following performance tools and debuggers have been updated in RHEL 9.3:

- **Valgrind 3.21**
- **SystemTap 4.9**
- **elfutils 0.189**

Updated performance monitoring tools

The following performance monitoring tools have been updated in RHEL 9.3:

- **PCP 6.0.5**
- **Grafana 9.2.10**

Updated compiler toolsets

The following compiler toolsets have been updated in RHEL 9.3:

- **GCC Toolset 13 (new)**

- **LLVM Toolset 16.0.6**
- **Rust Toolset 1.71.1**
- **Go Toolset 1.20.10**

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

Java implementations in RHEL 9

The RHEL 9 AppStream repository includes:

- The **java-21-openjdk** packages, which provide the OpenJDK 21 Java Runtime Environment and the OpenJDK 21 Java Software Development Kit. An OpenJDK 21.0.1 security release is also available to install. It is recommended that you install the OpenJDK 21.0.1 update to acquire the latest security fixes.
- 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 9.3 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](#).

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, RHEL 8.8 to RHEL 9.2, and RHEL 8.9 to RHEL 9.3 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 before 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:

- Requirements on disk space have been significantly reduced on systems with XFS filesystems formatted with **ftype=0**.
- Disk images created during the upgrade process for upgrade purposes now have dynamic sizes. The **LEAPP_OVL_SIZE** environment variable is not needed anymore.
- Issues with the calculation of the required free space on existing disk partitions have been fixed. The missing free disk space is now correctly detected before the required reboot of the system, and the report correctly displays file systems that do not have enough free space to proceed the upgrade RPM transaction.
- Third-party drivers can now be managed during the in-place upgrade process using custom leapp actors.
- An overview of the pre-upgrade and upgrade reports is now printed in the terminal.
- Upgrades of RHEL Real Time and RHEL Real Time for Network Functions Virtualization (NFV) in Red Hat OpenStack Platform are now supported.

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.



NOTE

Public release notes include links to access the original tracking tickets, but private release notes are not viewable so do not include links.^[1]

[1] Public release notes include links to access the original tracking tickets, but private release notes are not viewable so do not include links.

CHAPTER 2. ARCHITECTURES

Red Hat Enterprise Linux 9.3 is distributed with the kernel version 5.14.0-362.8.1, which provides support for the following architectures at the minimum required version (stated in parentheses):

- 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 operating system 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 9.3.

4.1. INSTALLER AND IMAGE CREATION

Support to both legacy and UEFI boot for AWS EC2 images

Previously, RHEL image builder created EC2 AMD or Intel 64-bit architecture AMIs images with support only for the legacy boot type. As a consequence, it was not possible to take advantage of certain AWS features requiring UEFI boot, such as secure boot. This enhancement extends the AWS EC2 AMD or Intel 64-bit architecture AMI image to support UEFI boot, in addition to the legacy BIOS boot. As a result, it is now possible to take advantage of AWS features which require booting the image with UEFI.

Jira:RHELDPCS-16339^[1]

New boot option `inst.wait_for_disks=` to add wait time for loading a Kickstart file or the kernel drivers

Sometimes, it may take a few seconds to load a Kickstart file or the kernel drivers from the device with the **OEMDRV** label during the boot process. To adjust the wait time, you can now use the new boot option, `inst.wait_for_disks=`. Using this option, you can specify how many seconds to wait before the installation. The default time is set to **5** seconds, however, you can use **0** seconds to minimize the delay. For more information about this option, see [Storage boot options](#).

Bugzilla:2171811

Ability to select required kernel while installing RHEL on ARM using GUI and TUI

Previously, you could install RHEL on ARM with kernel-64k page size only by using the Kickstart method. With this update, you can now install RHEL on ARM using the GUI or the TUI and selecting the required kernel version. The option to select the required kernel is available on the Software Selection screen under Kernel Options.

Bugzilla:2164819^[1]

Support for VMware VSphere (OVA)

This update adds support to build VMware VSphere OVA files by using RHEL image builder. The Open Virtual Appliance (OVA) file is a virtual appliance used by the VMware VSphere virtualization application. The OVA file contains files used to describe a virtual machine, such as an OVF descriptor file, one or more virtual machine disk image files (VMDK), optional manifest (MF) and certificate files. By using the VMware VSphere (.ova), you can more easily deploy the image to VMware vSphere by using the vSphere GUI client. You can further customize the resulting VM before you boot the image.

Jira:RHELDPCS-16877^[1]

New network Kickstart options to control DNS handling

You can now control DNS handling using the **network** Kickstart command with the following new options. Use these new options with the **--device** option.

- The **--ipv4-dns-search** and **--ipv6-dns-search** options allow you to set DNS search domains manually. These options mirror their NetworkManager properties, for example:

```
network --device ens3 --ipv4-dns-search domain1.example.com, domain2.example.com
```

- The **--ipv4-ignore-auto-dns** and **--ipv6-ignore-auto-dns** options allow you to ignore DNS settings from DHCP. They do not require any arguments.

Bugzilla:2065754^[1]

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^[1]

4.2. SECURITY

Keylime rebased to version 7.3.0

The Keylime packages have been updated to upstream version 7.3.0. This version provides various enhancements and bug fixes. Most notably, the allow and exclude lists are combined into the Keylime runtime policy. You can combine the two lists by using the **convert_runtime_policy.py** script.

In addition, the update fixes two vulnerabilities with the moderate impact rating: [CVE-2023-38200](#) and [CVE-2023-38201](#).

Jira:RHEL-476^[1]

Ports for Keylime have stricter rules in SELinux policy

Ports used by Keylime are now labeled as **keylime_port_t** in the Keylime SELinux policy. The policy now allows TCP connections for ports with this label. This is because the previous Keylime SELinux policy allowed connecting to all undefined ports and also most of the ports used by Keylime were in the undefined group. As a result, this update increases the granularity of the Keylime SELinux policy, and port security can be more strict and better targeted.

Jira:RHEL-595^[1]

Audit now supports **FANOTIFY** record fields

This update of the **audit** packages introduces support for **FANOTIFY** Audit record fields. The Audit subsystem now logs additional information in the **AUDIT_FANOTIFY** record, notably:

- **fan_type** to specify the type of a **FANOTIFY** event
- **fan_info** to specify additional context information
- **sub_trust** and **obj_trust** to indicate trust levels for a subject and an object involved in an event

As a result, you can better understand why the Audit system denied access in certain cases. This can help you write policies for tools such as the **fapolicyd** framework.

Jira:RHELPLAN-161087^[1]

fapolicyd now provides rule numbers for troubleshooting

With this enhancement, new kernel and Audit components allow the **fapolicyd** service to send the number of the rule that causes a denial to the **fanotify** API. As a result, you can troubleshoot problems related to **fapolicyd** more precisely.

[Jira:RHEL-624](#)

crypto-policies now provides the **NO-ENFORCE-EMS** subpolicy for TLS 1.2 connections in FIPS mode

The system-wide cryptographic policies now contain the **NO-ENFORCE-EMS** subpolicy. After applying the new subpolicy, the system no longer requires the Extended Master Secret (EMS) extension (RFC 7627) for all TLS 1.2 connections negotiated in FIPS mode. This allows the system to connect to legacy systems without support for EMS or TLS 1.3. Note that this violates the requirements of the FIPS-140-3 standard. You can apply the subpolicy by entering the **update-crypto-policies --set FIPS:NO-ENFORCE-EMS** command.

[Bugzilla:2216257^{\[1\]}](#)

GnuTLS requires EMS with TLS 1.2 in FIPS mode

To comply with the FIPS-140-3 standard, GnuTLS servers and clients require the Extended Master Secret (EMS) extension (RFC 7627) for all TLS 1.2 connections negotiated in FIPS mode. If your scenario requires preserving compatibility with older servers and clients that do not support EMS and you cannot use TLS 1.3, you can apply the **NO-ENFORCE-EMS** system-wide cryptographic subpolicy:

```
# update-crypto-policies --set FIPS:NO-ENFORCE-EMS
```



WARNING

If you allow TLS 1.2 connections without EMS, your system no longer meets the FIPS-140-3 requirements.

[Bugzilla:2157953](#)

NSS now enforce EMS in FIPS mode

The Network Security Services (NSS) libraries now contain the **TLS-REQUIRE-EMS** policy to require the Extended Master Secret (EMS) extension (RFC 7627) for all TLS 1.2 connections as mandated by the FIPS 140-3 standard. NSS use the new policy when the system-wide cryptographic policies are set to **FIPS**.

If your scenario requires interoperating with legacy systems without support for EMS or TLS 1.3, you can apply the **NO-ENFORCE-EMS** system-wide cryptographic subpolicy. Such a change violates the FIPS-140-3 requirements.

[Bugzilla:2157950](#)

OpenSSL now supports disabling EMS in FIPS mode

You can now configure the OpenSSL cryptographic libraries to allow for TLS 1.2 connections without the Extended Master Secret (EMS) extension (RFC 7627) in FIPS mode by editing the **/etc/pki/tls/fips_local.cnf** file. In a text editor of your choice, add the following section to the

configuration file:

```
[fips_sect]
tls1-prf-ems-check = 0
activate = 1
```

Then, locate the SSL configuration section in the `/etc/pki/tls/openssl.cnf` file. The default SSL configuration section is `crypto_policy`. At the end of the SSL configuration section, add the following line:

```
Options=RHNoEnforceEMSinFIPS
```

The previous configuration changes allow the system in FIPS mode to connect to legacy systems without support for EMS or TLS 1.3.



WARNING

You can stop enforcing EMS for TLS 1.2 in FIPS mode by entering the `update-crypto-policies --set FIPS:NO-ENFORCE-EMS` command. In both cases, such a configuration change violates the requirements of the FIPS-140-3 standard.

Bugzilla:2216256^[1]

OpenSSH further enforces SHA-2

As part of the effort to migrate further from the less secure SHA-1 message digest for cryptographic purposes, the following changes were made in OpenSSH:

- Added a check on `sshd` startup whether using SHA-1 is configured on the system. If it is not available, OpenSSH does not try to use SHA-1 for operations. This eliminates loading DSS keys when they are present and also enforces advertising `rsa-sha2` combinations when they are available.
- On SSH private key conversion, OpenSSH explicitly uses SHA-2 for testing RSA keys.
- When SHA-1 signatures are unavailable on the server side, `sshd` uses SHA-2 to confirm host key proof. This might be incompatible with clients on RHEL 8 and earlier versions.
- When the SHA-1 algorithm is unavailable on the client side, OpenSSH uses SHA-2.
- On the client side, OpenSSH permits SHA-2-based key proofs from the server when SHA-1 was used in key proof request or when the hash algorithm is not specified (assuming default). This is aligned with the already present exception for RSA certificates, and allows connecting by using modern algorithms when supported.

Bugzilla:2070163

OpenSSL now contains protections against Bleichenbacher-like attacks

This release of the OpenSSL TLS toolkit introduces API-level protections against Bleichenbacher-like attacks on the RSA PKCS #1 v1.5 decryption process. The RSA decryption now returns a randomly

generated deterministic message instead of an error if it detects an error when checking padding during a PKCS #1 v1.5 decryption. The change provides general protection against vulnerabilities such as [CVE-2020-25659](#) and [CVE-2020-25657](#).

You can disable this protection by calling the `EVP_PKEY_CTX_ctrl_str(ctx, "rsa_pkcs1_implicit_rejection", "0")` function on the RSA decryption context, but this makes your system more vulnerable.

[Bugzilla:2153471](#)

OpenSSL now supports Brainpool curves configurable through the `Groups` option

This update of the OpenSSL TLS toolkit introduces support for Brainpool curves in Elliptic Curve Cryptography (ECC). Additionally, you can control the curves with the system-wide cryptographic policies through the `Groups` configuration option.

The following Brainpool curves are enabled in OpenSSL ECC:

- `brainpoolP256r1`
- `brainpoolP256t1`
- `brainpoolP320r1`
- `brainpoolP320t1`
- `brainpoolP384r1`
- `brainpoolP384t1`
- `brainpoolP512r1`
- `brainpoolP512t1`

[Bugzilla:2188180](#)

`crypto-policies` now supports OpenSSL ECC Brainpool curves

With this update of the system-wide cryptographic policies, you can now control the following Brainpool Elliptic Curve Cryptography (ECC) curves in OpenSSL by using the `group` option:

- `BRAINPOOL-P256R1`
- `BRAINPOOL-P384R1`
- `BRAINPOOL-P512R1`.

For example, you can enable all supported Brainpool elliptic curves in OpenSSL by creating a subpolicy that contains the following line:

```
group = BRAINPOOL-^+ 
```

[Bugzilla:2193324^{\[1\]}](#)

`crypto-policies` now use the same group order as OpenSSL by default

In this release, the system-wide cryptographic policies (`crypto-policies`) control the group order in the OpenSSL `Groups` configuration option. To preserve the performance in OpenSSL, `crypto-policies` use

the default group order that matches the order of the OpenSSL built-in preferences. As a result, the RHEL cryptographic back ends that support **crypto-policies** for controlling the group order, such as GnuTLS, now use the same order as OpenSSL.

[Jira:RHEL-591^{\[1\]}](#)

crypto-policies permitted_encyptypes no longer break replications in FIPS mode

Before this update, an IdM server running on RHEL 8 sent an AES-256-HMAC-SHA-1-encrypted service ticket that an IdM replica running RHEL 9 in FIPS mode. Consequently, the default **permitted_encyptypes krb5** configuration broke a replication between the RHEL 8 IdM server and the RHEL 9 IdM replica in FIPS mode.

This update of the system-wide cryptographic policies reorders the **permitted_encyptypes krb5** configuration option values to allow prioritization of more interoperable encryption types by default. As a result, the **permitted_encyptypes** configuration no longer break replications between a RHEL 8 IdM servers and a RHEL 9 IdM replica in FIPS mode.



NOTE

If you use Kerberos, verify the order of the values of **permitted_encyptypes** in the **/etc/crypto-policies/back-ends/krb5.config** file. If your scenario requires a different order, apply a custom cryptographic subpolicy.

[Bugzilla:2225222](#)

pcsc-lite-ccid rebased to 1.5.2

The **pcsc-lite-ccid** package has been updated to version 1.5.2. This version provides various bug fixes and enhancements, most notably:

- Support for new readers
- Fix for Alcor Micro AU9560

[Bugzilla:2209457](#)

opensc rebased to 0.23

The **opensc** packages have been updated to version 0.23. This version provides various bug fixes and enhancements, most notably:

- Added support for encryption and decryption using symmetric keys
- Added support for signing data with a length of more than 512 bytes
- Disabled old card driver support by default
- Removed support for old drivers MioCOS and JCOP

[Jira:RHEL-280^{\[1\]}](#)

setools rebased to 4.4.3

The **setools** packages have been updated to version 4.4.3. This version provides various bug fixes and enhancements, most notably:

- Fixed compilation with Cython 3.0.0
- Improved man pages
- Removed unused options in **sediff**, **sesearch**, and **apol**
- Added the **-r** option to **seinfoflow** command to get flows analysis into the source type
- Rules with no permissions are automatically rejected as an invalid policy

[Bugzilla:2231801](#), [Bugzilla:2184140](#)

Additional services confined in the SELinux policy

This update adds additional rules to the SELinux policy that confine the following **systemd** services:

- **qat**
- **systemd-pstore**
- **boothd**
- **fdo-manufacturing-server**
- **fdo-rendezvous-server**
- **fdo-client-linuxapp**
- **fdo-owner-onboarding-server**

As a result, these services do not run with the **unconfined_service_t** SELinux label anymore, and run successfully in SELinux enforcing mode.

[Bugzilla:2080443^{\[1\]}](#), [Bugzilla:2026795](#), [Bugzilla:2181565](#), [Bugzilla:2128833](#)

OpenSCAP rebased to 1.3.8

The OpenSCAP packages have been rebased to upstream version 1.3.8. This version provides various bug fixes and enhancements, most notably:

- Fixed **systemd** probes to not ignore some **systemd** units
- Added offline capabilities to the **shadow** OVAL probe
- Added offline capabilities to the **sysctl** OVAL probe
- Added **auristorfs** to the list of network filesystems
- Created a workaround for issues with tailoring files produced by the **autotailor** utility

[Bugzilla:2217442](#)

SCAP Security Guide rebased to version 0.1.69

The SCAP Security Guide (SSG) packages have been rebased to upstream version 0.1.69. This version provides various enhancements and bug fixes. Most notably, it introduces three new SCAP profiles for RHEL 9 aligned with the CCN-STIC-610A22 Guide issued by the National Cryptologic Center of Spain in October 2022:

Profile name	Profile ID	Policy version
CCN Red Hat Enterprise Linux 9 - Advanced	xccdf_org.ssgproject.content_profile_ccn_advanced	2022-10
CCN Red Hat Enterprise Linux 9 - Basic	xccdf_org.ssgproject.content_profile_ccn_basic	2022-10
CCN Red Hat Enterprise Linux 9 - Intermediate	xccdf_org.ssgproject.content_profile_ccn_intermediate	2022-10

[Bugzilla:2221697](#)

ANSSI-BP-028 security profiles updated to version 2.0

The following French National Agency for the Security of Information Systems (ANSSI) BP-028 in the SCAP Security Guide were updated to be aligned with version 2.0:

- ANSSI-BP-028 Minimal Level
- ANSSI-BP-028 Intermediary Level
- ANSSI-BP-028 Enhanced Level
- ANSSI-BP-028 High Level

[Bugzilla:2155790](#)

python3-greenlet-devel is now available in CRB

The **python3-greenlet-devel** package is now available in the CodeReady Linux Builder (CRB) repository, which you must explicitly enable. See the [How to enable and make use of content within CodeReady Linux Builder](#) Knowledgebase article for more information. Note that packages included in the CRB repository are unsupported.

[Bugzilla:2149497](#)

SSG rule to check the group used by the pam_wheel.so module is simplified

The CIS Benchmark requires restricting the **su** command in favor of the **sudo** command. SCAP Security Guide (SSG) fulfills this requirement with the **pam_wheel.so** module, which restricts the **su** command to a specific group. This update improves the rule that checks whether this group exists and has no members. As a result, the rule is more efficient and simplifies the interpretation of the assessment report.

[Jira:RHEL-1905](#)

4.3. RHEL FOR EDGE

New FIDO Device Onboarding Servers container images are available

The following FIDO Device Onboarding Servers container images for onboarding IoT and edge computing devices are now available in the [Red Hat Container Catalog](#):

- rhel9/fdo-manufacturing-server container image

- `rhel9/fdo-owner-onboarding-server` container image
- `rhel9/fdo-rendezvous-server` container image
- `rhel9/fdo-serviceinfo-api-server` container image

[Jira:RHELPLAN-163133^{\[1\]}](#)

The `minimal-raw` image type now supports 64-bit ARM architectures

With this enhancement, you can create a **minimal-raw** image type with support for 64-bit ARM architecture, and AMD and Intel 64-bit architectures. The **minimal-raw** image is pre-packaged, bootable, minimal RPM image, compressed in the **xz** format. To boot the image, you must decompress it and copy to any bootable device, such as an SD card. To decompress the image, run the following command:

```
$ xz -d <_uuid-minimal-raw.img_.xz>
```

[Jira:RHELPLAN-163665^{\[1\]}](#)

The Commit ID is now supported as a value for the `--parent` argument of `composer-cli` CLI

You can now use the image Commit ID as a value for the `--parent` argument of the `composer-cli` command line. To get the image Commit ID, download and extract the RHEL for Edge Commit image. You can find the **ref** name and the commit ID in the extracted **.tar** file.

[Jira:RHELDPCS-16386^{\[1\]}](#)

Support to build RHEL for Edge `.ami` images

With this enhancement, you have support to build **.ami** images for RHEL for Edge by using on-premise RHEL image builder. During the initial boot, you can customize the blueprint with Ignition to inject the credentials into the image. You can upload the **.ami** image to AWS and boot an EC2 instance in AWS.

[Jira:RHELDPCS-16708^{\[1\]}](#)

Support to build `.vmdk` images for RHEL for Edge

With this enhancement, you have support to build a **.vmdk** image for RHEL for Edge by using on-premise RHEL image builder. You can customize the blueprint with Ignition to inject the credentials into the image during the initial boot. You can load the image on vSphere and boot the image in a VM vSphere. The image is compatible with ESXi 7.0 U2, ESXi 8.0, and later. The VM is compatible with versions 19 and 20.

[Jira:RHELDPCS-16709^{\[1\]}](#)

You can now log in to an Edge system as the initial user without setting a password

Previously, logging in as the initial user created during the FDO onboarding process did not work because the system asked for a password that was not set with the **useradd** command. With this enhancement, the password is now set to optional, and you can log in even if you did not previously set a password by using the **useradd** command. Note that you can log in with an SSH key without entering a password, and if it fails, you will be prompted to enter a password.

[Jira:RHELDPCS-17101^{\[1\]}](#)

4.4. SOFTWARE MANAGEMENT

New DNF Automatic **reboot** option for an automatic reboot after an upgrade

With this enhancement, you can use the DNF Automatic **reboot** option to set your system to automatically reboot to apply the changes after an upgrade.

The **reboot** option supports the following settings:

- **never** does not reboot the system. This is the current behavior.
- **when-changed** triggers a reboot after any upgrade.
- **when-needed** triggers a reboot only when rebooting is required to apply changes, for example, when systemd or the kernel is upgraded.

You can use the **reboot_command** option to customize the command used to reboot. The default reboot command is **shutdown -r**.

[Bugzilla:2124793](#)

The new **--poweroff** option allows you to shut down the system after installing updates

With this enhancement, the new **--poweroff** option has been added to the **reboot** command of the **dnf system-upgrade** plugin. You can use this option to shut down the system after installing updates instead of rebooting.

[Bugzilla:2157844](#)

New **dnf leaves** and **show-leaves** plug-ins are now available for the DNF API

With this enhancement, the following new DNF plug-ins are available that list packages installed on your system that are not required as dependencies of other installed packages:

- **dnf leaves** lists all packages.
- **show-leaves** lists newly installed packages and packages that became unrequired as dependencies of other installed packages after a transaction.

[Bugzilla:2134638](#)

4.5. SHELLS AND COMMAND-LINE TOOLS

The NetBackup services are now enabled for backup restoration

When using the NetBackup (NBU) backup method, ReaR now includes the unit files for the NetBackup services version 10.1.1 in the rescue image and starts them when the rescue system boots. As a result, you can restore the system backup by using the NBU backup method during the recovery process and complete the restore successfully.

[Bugzilla:2188593](#)

opencryptoki rebased to 3.21.0

The **opencryptoki** package has been rebased to version 3.21.0, which provides many enhancements and bug fixes. Most notably, **opencryptoki** now supports the following features:

- Concurrent hardware security module (HSM) master key changes

- The **protected-key** option to transform a chosen key into a protected key
- Additional key types, such as DH, DSA, and generic secret key types
- EP11 host library version 4
- AES-XTS key type
- IBM-specific Kyber key type and mechanism
- Additional IBM-specific Dilithium key round 2 and 3 variants

Additionally, **pkcs11slotd** slot manager no longer runs as root and **opencryptoki** offers further hardening. With this update, you can also use the following set of new commands:

p11sak set-key-attr

To modify keys

p11sak copy-key

To copy keys

p11sak import-key

To import keys

p11sak export-key

To export keys

[Bugzilla:2160061^{\[1\]}](#)

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 IPoIB devices using systemd link file](#) procedure.

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

[Bugzilla:2136937](#)

4.6. INFRASTRUCTURE SERVICES

Postfix now supports SRV lookups

With this enhancement, you can now use the Postfix DNS service records resolution (SRV) to automatically configure mail clients and balance load of servers. Additionally, you can prevent mail delivery disruptions caused by temporary DNS issues or misconfigured SRV records by using the following SRV-related options in your Postfix configuration:

use_srv_lookup

You can enable discovery for the specified service by using DNS SRV records.

allow_srv_lookup_fallback

You can use a cascading approach to locating a service.

ignore_srv_lookup_error

You can ensure that the service discovery remains functional even if SRV records are not available or encounter errors.

[Bugzilla:2134789](#)

Generic LF-to-CRLF driver is available in cups-filters

With this enhancement, you can now use the Generic LF-to-CRLF driver, which converts LF characters to CR+LF characters for printers accepting files with CR+LF characters. The carriage return (CR) and line feed (LF) are control characters that mark the end of lines. As a result, by using this driver, you can send an LF character terminated file from your application to a printer accepting only CR+LF characters. The Generic LF-to-CRLF driver is a renamed version of the **text-only** driver from RHEL 7. The new name reflects its actual functionality.

[Bugzilla:2229784](#)

4.7. NETWORKING

RHEL on ARM now fully supports wifi adapters in RHEL 9.3

With this enhancement, you can now enable access to wifi adapters for several cards for the **arm64** platforms.

For details on configuring wifi connections, see [Managing wifi connections](#).

[Bugzilla:2208365^{\[1\]}](#)

NetworkManager now supports the no-aaaa option in resolv.conf

NetworkManager now supports adding the **no-aaaa** DNS option in the resolv.conf file. By using the **no-aaaa** value in the DNS option setting, you can disable IPv6 DNS resolution.

[Bugzilla:2176137](#)

nmstate now supports mixing static DNS search along with dynamic DNS name servers

The **nmstate** framework now supports both static Domain Name System (DNS) search domains and dynamic DNS name servers, which **nmstate** obtained from Dynamic Host Configuration Protocol (DHCP) or the **autoconf** mechanism. Previously, static DNS search domains could not co-exist with dynamic DNS name servers because the dynamic configurations were discarded by **nmstate**. This often led to unnecessary complexity and limitations in network setup and management. This enhancement aims to bring more flexibility in managing DNS configurations. As a result, **nmstate** attempts to find a network interface to store the DNS configuration in the following order:

1. The preferred interface, which currently holds the DNS configuration and is still valid for DNS
2. An automatic interface
3. An IP enabled interface

Note that this enhancement does not remove the DNS name servers learned from DHCP.

The following is an example YAML file to apply this feature:

```
---
dns-resolver:
  config:
```

```
search:
  - example.com
  - example.org
interfaces:
  - name: eth1
    type: ethernet
    state: up
  ipv4:
    enabled: true
    dhcp: true
  ipv6:
    enabled: true
    dhcp: true
    autoconf: true
```

[Bugzilla:2179916](#)

nmstate now supports the `bridge.vlan-default-pvid` NetworkManager configuration option

With this update, you can use the **nmstate** framework to configure the **bridge.vlan-default-pvid** NetworkManager configuration option. By using this option, you can set the default port VLAN identifier (PVID) for untagged traffic on a bridge interface that supports VLANs, when you use Linux bridge VLAN filtering. To achieve this result, use the following YAML configuration:

```
interfaces:
  - name: linux-br0
    type: linux-bridge
    state: up
    bridge:
      options:
        vlan-default-pvid: 5
      port:
        - name: eth1
          stp-hairpin-mode: false
          stp-path-cost: 100
          stp-priority: 32
        vlan:
          mode: access
          tag: 100
```

Note that the default value of **bridge.vlan-default-pvid** is 1. When set to 0 with VLAN filtering enabled, the untagged traffic is dropped.

[Bugzilla:2180795](#)

The NetworkManager service restarts immediately after the `dbus` service is restarted

Previously, after restarting **dbus** for some reason, **NetworkManager** stopped. This behavior was not optimal and caused a loss of connectivity. Therefore, this enhancement updates **NetworkManager** to become more robust and to make it restart automatically upon a **dbus** restart.

[Bugzilla:2161915](#)

The `nm-cloud-setup` utility now supports IMDSv2 configuration

Users can configure an AWS Red Hat Enterprise Linux EC2 instance with Instance Metadata Service Version 2 (IMDSv2) with the **nm-cloud-setup** utility. To comply with improved security that restricts

unauthorized access to EC2 metadata and new features, integration between AWS and Red Hat services is necessary to provide advanced features. This enhancement enables the **nm-cloud-setup** utility to fetch and save the IMDSv2 tokens, verify an EC2 environment, and retrieve information about available interfaces and IP configuration by using the secured IMDSv2 tokens.

[Bugzilla:2151986](#)

NetworkManager notifies when using the deprecated **ifcfg** format

Connection profiles in **ifcfg** format are deprecated in RHEL 9 (see [NetworkManager connection profiles in ifcfg format are deprecated](#)). With this update, NetworkManager notifies users about the deprecation of this format:

- NetworkManager logs the following warning to the **systemd** journal if it processes a connection profile in **ifcfg** format in the `/etc/sysconfig/network-scripts/` directory:

```
Warning: the ifcfg-rh plugin is deprecated, please migrate connections to the keyfile format using "nmcli connection migrate"
```

- The **nmcli** utility reports the following error if you try to modify a property that is not supported in **ifcfg** format:

```
Error: Failed to modify connection '<name>': failed to update connection: The ifcfg-rh plugin doesn't support setting '<property>'. If you are modifying an existing connection profile saved in ifcfg-rh format, please migrate the connection to keyfile using 'nmcli connection migrate <connection_uuid>' or via the Update2() D-Bus API and try again.
```

As a result of these enhancements, NetworkManager now notifies users if they still use or modify connection profiles in the deprecated **ifcfg** format.

For further details about migrating profiles from **ifcfg** to keyfile format, see [Migrating NetworkManager profiles from ifcfg to keyfile format](#).

[Bugzilla:2190375](#)

NetworkManager now supports the **lACP_active** option in the bonding configuration

By using **NetworkManager**, the **lACP_active** option in bonding configuration provides fine-grained control over Link Aggregation Control Protocol Data Units (LACPDU) frames. The **lACP_active** option also adjusts the behavior of LACPDU frames and controls periodic transmission of these frames in the bonding setup. To customize network configurations, you can enable or disable periodic transmission of LACPDU frames by setting **lACP_active** to **ON** or **OFF**.

[Bugzilla:2069001](#)

NetworkManager now supports configuration of the **ns_ip6_target** option for bond interfaces

This enhancement allows setting the **arp_interval** option by specifying a maximum of 16 IPv6 addresses as monitoring peers in **NetworkManager** for configuration of the **ns_i6_target** option for bond interfaces. Previously, it was not possible to specify IPv6 monitoring peers in **NetworkManager**. With this update, you can configure the **ns_ip6_target** option in the **bond.options** parameter by using the **nmcli** utility. **NetworkManager** applies this setting to the bond interface by enabling the specification of a maximum of 16 IPv6 addresses. This enhancement equally applies to IPv4 and IPv6 settings.

[Bugzilla:2069004](#)

NetworkManager now supports both static and DHCP IP configuration on the same network interface

By using the **nmstate** utility, you can now assign a static IP address with **dhcp: true** or **autoconf: true** value on the DHCP or Ad-Hoc Network Autoconfiguration (autoconf) enabled interface.

With this enhancement, **nmstate** supports two properties of IP addresses:

- **valid_lft** means valid lifetime in seconds
- **preferred_lft** means preferred lifetime in seconds

Default value of both parameters is **forever** which means static.

With above properties, **nmstate** can ignore DHCP/autoconf based IP addresses to avoid converting dynamic IP addresses to static IP after applying the queried state back. If your scenario requires having disabled DHCP/autoconf settings with dynamic IP addresses, **nmstate** converts those dynamic IP to static IP addresses.

[Bugzilla:2177733](#)

nmstate supports MAC address identifiable network interface

The **nmstate** utility supports network configuration directly to a network interface with a MAC address instead of an interface name.

This enhancement introduces two properties to the base interface:

- **identifier** : identifies **name** or **mac-address** on a network. The default value is **name**.
- **profile-name** : string

When the **identifier** variable is set to the **mac-address** value, **nmstate** uses the **interface.mac-address** over **interface.name** to choose a network interface for a specific network state. When storing the network configuration, if the **interface.profile-name** variable is not assigned, **nmstate** prefers **interface.profile-name** over **interface.name**. If you check the current network state, the **interface.profile-name** remains hidden if it is equal to **interface.name**.

[Bugzilla:2183214](#)

NetworkManager supports defining after how many failed ARP checks the bonding driver marks a port as down

This enhancement adds the **arp_missed_max** option to bond connection profiles in NetworkManager. If you use the Address Resolution Protocol (ARP) monitor to check if ports of a bond are up, you can now set **arp_missed_max** to define after how many failed checks the bonding driver marks the port as down.

[Bugzilla:2148684](#)

NetworkManager supports specifying link-related properties

This enhancement adds the following network link properties to NetworkManager connection profiles:

- **link.tx-queue-length** - The size of the transmit (TX) queue length in number of packets.
- **link.gro-max-size** - The maximum size in bytes of a Generic Receive Offload (GRO) packet the device accepts.

- **link.gso-max-segments** - The maximum number of segments of a Generic Segmentation Offload (GSO) packet the device accepts.
- **link.gso-max-size** - The maximum size in bytes of a GSO packet.

Previously, you could configure these kernel settings only by using **ip** commands or by using such commands in NetworkManager dispatcher scripts. With this enhancement, you can now configure these settings directly in connection profiles.

Note that NetworkManager supports these properties only in connection profiles in **keyfile** format and not in the deprecated **ifcfg** format.

[Bugzilla:2158328](#)

The nmstate API support available for **dhcp-send-hostname** and **dhcp-custom-hostname** DHCP options

With this enhancement, the **nmstate** utility supports configuration of the following two DHCP options in the connection file:

- **dhcp-send-hostname: true** or **false** value. If a DHCP request needs the hostname or fully qualified domain name (FQDN) option, the hostname from that option is set. The default is **true**.
- **dhcp-custom-hostname: <string>**. Use this option to configure the hostname or FQDN option in a DHCP request, value type is string.

For DHCPv4 network protocols

- If the hostname is FQDN, see the **Fully Qualified Domain Name (FQDN)**, option (81) in RFC 4702.
- If the hostname is not FQDN, see the **Host Name**, option (12) in RFC 2132.

For DHCPv6 network protocols

Supports custom string, empty domain name, overrides the hostname for a DHCP request. See the **Fully Qualified Domain Name (FQDN)**, option (29) in RFC 4704.

[Bugzilla:2187622](#)

NetworkManager rebased to version 1.44.0

The **NetworkManager** packages have been upgraded to upstream version 1.44.0, which provides several enhancements and bug fixes over the previous version:

- [Link-related properties have been added to NetworkManager](#) .
- The **arp_missed_max**, **lACP_active**, and **ns_ip6_target** properties have been added to bond connection profiles.
- You can now set a DHCPv6 prefix delegation hint in the **ipv6.dhcp-pd-hint** connection property.
- Enabling the new **rename** parameter in the **[keyfile]** section of the **/etc/NetworkManager/NetworkManager.conf** file causes NetworkManager to rename a connection profile in **/etc/NetworkManager/system-connections/** if you change a profile name

(**connection.id**). If external applications or scripts rely on the file names, do not enable this parameter.

- When you set a hostname that contains a non-public top-level domain (TLD), NetworkManager now uses this TLD as DNS search domain instead of the full hostname.
- NetworkManager now applies DNS options from the **[global-dns]** section in the **/etc/NetworkManager/NetworkManager.conf** file.
- To avoid race conditions with other depending services, NetworkManager now acquires the D-Bus name only after populating the D-Bus tree. Note that this can add a delay when NetworkManager starts.
- NetworkManager now adds a **version-id** argument to **Update2()** D-Bus calls to prevent concurrent profile modifications.
- NetworkManager no longer uses tentative IPv6 addresses to resolve the system hostname from DNS.
- To prevent unexpected behaviors in case of multi-connect profiles, NetworkManager now tracks the number of auto-connect retries left for each device and connection instead of only per connection.
- NetworkManager sets VLAN filtering options by using the kernel's **netlink** interface instead of the **sysfs** file system.
- The **nm-cloud-setup** utility now supports Instance Metadata Service Version 2 (IDMSv2) on Amazon EC2.
- Users can now enable and disable wifi and Wireless Wide Area Networks (WWANs) in the **nmtui** application.
- Bond, bridge, and team connections now use the **ignore-carrier=no** setting in the **[main]** section of the **/etc/NetworkManager/NetworkManager.conf** file.

[Bugzilla:2180966](#)

SCTP rebased to the latest version of the kernel networking tree for RHEL 9

Notable changes in the Stream Control Transmission Protocol (SCTP) networking subsystem include:

- Virtual routing and forwarding (VRF) support to segment and isolate SCTP traffic within complex network environments.
- New stream schedulers (**fair capacity**, and **weighted fair queueing**) to ensure efficient and equal resource allocation in the network.

[Bugzilla:2189292](#)

MPTCP rebased to the latest version of the kernel networking tree for RHEL 9

Notable changes in the Multipath TCP (MPTCP) protocol extension include:

- Support for TCP fastopen (TFO) extension, including the client-side support. This feature offers latency, efficiency, and performance improvements for your network.
- Support multiple mixed IPv4/IPv6 subflows to allow for greater flexibility and adaptability in networks where both IP versions are used.

Bugzilla:2193330^[1]

The **xdp-tools** package rebased to version 1.4.0

The **xdp-tools** package has been upgraded to version 1.4.0, which provides multiple bug fixes and enhancements. Notable changes include:

- The **xdp-bench** utility gained support for multi-buffer eXpress Data Path (XDP) and for benchmarking the **xdp_load_bytes()** helper in the kernel. This feature enables conducting network benchmarking tests with large maximum transmission units (MTUs).
- The locking of the command line utilities of **xdp-tools** was improved to prevent stale locks if the utility did not exit cleanly.
- The **libxdp** library contains a new **xsk_umem__create_with_fd()** API that accepts an extra file descriptor of an already open **AF_XDP** socket. You can use this function as a substitute for the regular **xsk_umem__create()** function when a process does not have **CAP_NET_RAW** privileges.

Bugzilla:2218500

iproute rebased to version 6.2.0

The **iproute** packages have been upgraded to upstream version 6.2.0, which provides several enhancements and bug fixes over the previous version. The most notable changes are:

- The new **ip stats** command manages and shows interface statistics. By default, the **ip stats show** command displays statistics for all network devices, including bridges and bonds. You can filter the output by using the **dev** and **group** options. For further details, see the **ip-stats(8)** man page.
- The **ss** utility now provides the **-T (--threads)** option to display thread information, which extends the **-p (--processes)** option. For further details, see the **ss(8)** man page.
- You can use the new **bridge fdb flush** command to remove specific forwarding database (fdb) entries which match a supplied option. For further details, see the **bridge(8)** man page.

Jira:RHEL-428^[1]

The kernel supports activating bond ports in a specific order

With this enhancement, the kernel's **netlink** interface supports setting a priority on each port if you configure a bond in **active-backup**, **balance-tlb** or **balance-alb** mode. The priority value uses a 32-bit Integer, and a higher value means a higher priority. As a result, you can now activate the bond ports in a specific order.

To use this feature, you can configure the priority by setting the **bond-port.prio** property when you create or modify a NetworkManager port connection profile.

Bugzilla:2092194^[1]

firewalld now avoids unnecessary firewall rule flushes

With the release of the [RHBA-2023:7748](#), advisory the **firewalld** service was upgraded in a sense that it will not remove all the existing rules from the **iptables** configuration if both following conditions are met:

- **firewalld** is using the **nftables** backend.

- There are no firewall rules created with the **--direct** option.

This change aims at reducing unnecessary operations (firewall rules flushes) and improves integration with other software.

[Jira:RHEL-14694^{\[1\]}](#)

Introduction of new `nmstate` attributes for the VLAN interface

With this update of the `nmstate` framework, the following VLAN attributes were introduced:

- **registration-protocol**: VLAN Registration Protocol. The valid values are **gvrp** (GARP VLAN Registration Protocol), **mvrp** (Multiple VLAN Registration Protocol), and **none**.
- **reorder-headers**: reordering of output packet headers. The valid values are **true** and **false**.
- **loose-binding**: loose binding of the interface to the operating state of its primary device. The valid values are **true** and **false**.

Your YAML configuration file can look similar to the following example:

```
---
interfaces:
- name: eth1.101
  type: vlan
  state: up
  vlan:
    base-iface: eth1
    id: 101
    registration-protocol: mvrp
    loose-binding: true
    reorder-headers: true
```

[Jira:RHEL-19142^{\[1\]}](#)

4.8. KERNEL

Kernel version in RHEL 9.3

Red Hat Enterprise Linux 9.3 is distributed with the kernel version 5.14.0-362.8.1.

[Bugzilla:2232554](#)

Support added for NVIDIA Grace CPUs

Red Hat Enterprise Linux 9.3 adds support for the NVIDIA Grace ARM 64-bit CPU.

[Jira:RHELDPCS-17055^{\[1\]}](#)

The RHEL kernel now supports AutoIBRS

Automatic Indirect Branch Restricted Speculation (AutoIBRS) is a feature provided by the AMD EPYC 9004 Genoa family of processors and later CPU versions. AutoIBRS is the default mitigation for the Spectre v2 CPU vulnerability, which boosts performance and improves scalability.

[Bugzilla:1898184^{\[1\]}](#)

perf rebased to version 6.2

The **perf** performance analysis tool has been rebased to version 6.2. Apart from numerous minor bug fixes and updates, the **perf list** command now displays Performance Monitor Unit (PMU) events that contain human-friendly names and descriptions. In addition, this update adds support for the following processors:

- Intel 13th generation of Core processors (Intel Raptor Lake-S)
- Intel 14th generation of processors (Intel Meteor Lake)
- Intel 5th generation Xeon server processors (Intel Emerald Rapids)

Bugzilla:2177180^[1]

The Intel® QAT kernel driver rebased to upstream version 6.2

The Intel® Quick Assist Technology (QAT) has been rebased to upstream version 6.2. The Intel® QAT includes accelerators optimized for symmetric and asymmetric cryptography, compression performance, and other CPU intensive tasks.

The rebase includes many bug fixes and enhancements. The most notable enhancement is the support available for following hardware accelerator devices for QAT GEN4:

- Intel Quick Assist Technology 401xx devices
- Intel Quick Assist Technology 402xx devices

Bugzilla:2144528^[1]

vTPM functionality is available for Linux containers

This enhancement introduces virtual Trusted Platform Module (**vTPM**) for Linux containers and other virtual environments. **vTPM** is a virtualized version of TPM that provides a dedicated TPM instance to use for a secure running environment. With **vTPM** proxy drivers, programs interact with an emulated TPM the same way as they interact with physical TPMs.

As a result, each virtual machine can now have a dedicated **vTPM** instance that is isolated and encrypted.

Bugzilla:2210263^[1]

crash rebased to version 8.0.3

crash is an interactive utility to analyze a running system and a core dump file created by **kdump** in case of a kernel crash. The **crash** utility has been rebased to version 8.0.3 that includes many bug fixes and enhancements. The most notable enhancement is the added IPv6 support.

For network interfaces that support IPv6, **crash** prints IPv6 addresses with the **net** or **net -s** command.

- The **net** command displays the list of network devices, names, and the IP address.
- The **net -s** command displays the following information:
 - The open network socket and sock addresses
 - The family and the type of sockets and sock addresses

- The source and destination address and ports for **INET** and **INET6** families

[Bugzilla:2170283](#)

LVM thin-provisioned storage volumes supported as the **vmcore** dump target

The **kdump** mechanism now supports thin-provisioned logical volumes as the **vmcore** target. To configure LVM thin provisioning, complete the following steps:

1. Create an LVM volume group.

```
vgcreate vg00 /dev/sdb
```

2. Create an LVM thin pool of 10 MB available space.

```
lvcreate -L 10M -T vg00/thinpool
```

3. Create an LVM thin volume with 300 MB of the file system space.

```
lvcreate -V 300M -T vg00/thinpool -n thinvol  
mkfs.ext4 /dev/vg00/thinvol
```

4. Configure the LVM thin pool threshold to automatically extend the space.

```
cat /etc/lvm/lvm.conf  
activation {  
    thin_pool_autoextend_threshold = 70  
    thin_pool_autoextend_percent = 20  
    monitoring = 1  
}
```

5. Enable the LVM thin pool monitoring service for the first kernel.

```
systemctl enable lvm2-monitor.service  
systemctl start lvm2-monitor.service
```

6. Append the following lines to the **kdump.conf** file to set the LVM thin volume as the **kdump** target.

```
ext4 /dev/vg00/thinvol  
path /
```

7. Start the **kdump** service.

```
kdumpctl restart
```

8. Verify the configuration by triggering a kernel panic and check if the **vmcore** is saved to **/dev/vg00/thinvol**.

As a result, with this enhancement, the **kdump** mechanism now extends capability to save the **vmcore** dump files on thin-provisioned storage volumes.

[Bugzilla:2083475](#)

makedumpfile rebased to upstream version 1.7.3

The **makedumpfile** tool, which makes the crash dump file small by compressing pages or excluding memory pages that are not required, has been rebased to upstream version 1.7.3. The rebase includes many bug fixes and enhancements.

The most notable change is the added 5-level paging mode for standalone dump (**sadump**) mechanism on AMD and Intel 64-bit architecture. The 5-level paging mode extends the processor's linear address width to allow applications access larger amounts of memory. 5-level paging extends the size of virtual addresses from 48 to 57 bits and the physical addresses from 46 to 52 bits.

[Bugzilla:2173815](#)

Red Hat Enterprise Linux supports ARM's SystemReady ES and IR tier

Red Hat Enterprise Linux now supports ARM's SystemReady ES and IR, while previously only the SR tier was supported. In RHEL 9.3, the NVIDIA Orin, NXP i.MX 8M, and NXP i.MX 8M Mini modules have been enabled and are candidates for the RHEL hardware certification. Hardware partners are able to [submit certifications](#) by enrolling in the Red Hat hardware certification journey. Customers can use the supported hardware listed in the catalog for an improved experience in production.

[Bugzilla:2195986^{\[1\]}](#)

RHEL on ARM now supports Bluetooth

With this enhancement, you can configure a bluetooth device by using the **bluetoothctl** tool on the command-line interface.

[Bugzilla:2187856^{\[1\]}](#)

RHEL on ARM now fully supports USB-attached cameras in RHEL 9.3

This enhancement enables the **CONFIG_MEDIA_SUPPORT** kernel configuration for RHEL on AMD and Intel 64-bit architectures platforms. With that, you can now use USB cameras on AMD and Intel 64-bit architectures systems.

[Bugzilla:2192722^{\[1\]}](#)

bpf rebased to version 6.3

The Berkeley Packet Filter (BPF) facility has been rebased to Linux kernel version 6.3. Notable changes and enhancements include:

- BPF trampoline is now available on the 64-bit IBM Z architecture.
- A new map type - **BPF_MAP_TYPE_USER_RINGBUF** - and related helpers have been defined for the communication between the user space and kernel over a BPF-specific ring buffer.
- BPF now provides new complex data structures: linked list and **rbtree**.
- BPF trampoline that traces programs now supports **struct** arguments.
- BPF now provides a way to export XDP features supported by a NIC.
- Hardware metadata are now exposed to XDP programs by using the BPF kernel functions (**kfuncs**) with initial support for RX hash and timestamp metadata.

- BPF now provides a helper that sets source and destination NAT addresses and ports in new **contrack** module entries in BPF programs.
- BPF can now write directly to the **nf_conn:mark** connection mark of the netfilter packet filtering framework.

[Bugzilla:2178930^{\[1\]}](#)

4.9. BOOT LOADER

New default behavior of **grub2-mkconfig** with BLS

In the Boot Loader Specification (BLS) framework, GRUB generates the boot menu dynamically from BLS snippets at boot, and it is not predefined in the **grub.cfg** file.

Previously, the **grub2-mkconfig** command generated a new **grub.cfg** file and always overwrote the command-line arguments in all BLS snippets with the value of the **GRUB_CMDLINE_LINUX** variable found in the **/etc/default/grub** file.

With this release, the **grub2-mkconfig** command no longer overwrites the kernel command line in BLS snippets with **GRUB_CMDLINE_LINUX** by default. Each kernel in the boot loader menu takes its kernel command line from its BLS snippet. This new default behavior is caused by the **GRUB_ENABLE_BLSCFG=true** option.

To regenerate **grub.cfg** so that kernels ignore BLS snippets and take their command line from **GRUB_CMDLINE_LINUX** instead, set the **GRUB_ENABLE_BLSCFG=false** option.

To update the kernel command line in BLS snippets according to **GRUB_CMDLINE_LINUX**, add the **--update-bls-cmdline** option:

```
# grub2-mkconfig -o /path/to/grub.cfg --update-bls-cmdline
```

Also note that you can make changes to BLS snippets for individual kernels using **grubby**:

```
# grubby --update-kernel /path/to/kernel --args "new args"
```

[Jira:RHELDOCS-16752^{\[1\]}](#)

4.10. FILE SYSTEMS AND STORAGE

NFS server now implements courteous server code for **nfsd**

This update introduces the implementation of courteous server code for **nfsd** in the RHEL kernel NFS server. With this new feature, the NFS server avoids revoking leases for clients that have lost contact with the server for an extended period, provided that there is no conflicting access while the client is out of contact.

[Bugzilla:2180124](#)

DAX mount option and reflink are now compatible

With this update, reflinked files are now generally compatible with DAX mode. The file system DAX mount option **-o dax=always** is compatible with reflink-enabled file systems. Files that were reflinked can be set to DAX mode using inode flags. For more information see the **xfs(5)** man page.

[Bugzilla:2192730^{\[1\]}](#)

New encryption types for the RPCSEC GSS Kerberos V5

The RPCSEC GSS Kerberos V5 mechanism now supports encryption types defined in RFC 6803 (Camellia Encryption for Kerberos 5) and RFC 8009 (AES Encryption with HMAC-SHA2 for Kerberos 5).

The following encryption types have been added:

- **camellia128-cts-cmac**
- **camellia256-cts-cmac**
- **aes128-cts-hmac-sha256-128**
- **aes256-cts-hmac-sha384-192**

This allows NFS clients and NFS servers to use stronger encryption types when negotiating GSS contexts.

[Bugzilla:2178741](#)

fuse3 now allows invalidating a directory entry without triggering **umount**

With this update, a new mechanism has been added to **fuse3** package, that allows invalidating a directory entry without automatically triggering the **umount** of any mounts that exists on the entry.

[Bugzilla:2188182](#)

Stratis storage manager is now available

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
- Pool Level Encryption
- TMP2 and NBDE Support

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

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

[Bugzilla:2041558](#)

Improvements to GFS2 file system configuration and operation

The following updates have been implemented for GFS2 file systems:

- The **mkfs.gfs2** command now supports the new **-U** option, which makes it possible to specify the file system UUID for the file system you create. If you omit this option, the file system's UUID is generated randomly.

- The **gfs2_jadd** command creates journals at a much faster speed than in previous releases.
- The GFS2 man pages have been improved.

[Bugzilla:2170017](#)

dmpd rebased to version 1.0.2

The **dmpd** package has been upgraded to version 1.0.2. Notable changes include:

- Rewriting the tools in the Rust language for memory safety and for using multiple threads to boost performance.
- Improving the **thin_check** and **cache_check** tools to save the time of LVM pool activation along with the system startup. The required execution time for these tools is now improved by more than ten times as compared to the previous version.
- Updating **thin_dump** and **thin_restore** tools to avoid losing sharing of the metadata **btrees** for snapshots. Now the restored metadata does not require more space.
- Adding new **thin_metadata_pack** and **thin_metadata_unpack** tools to compress thin metadata, typically to a tenth of the size. This is better than the generic compressors. With this tool, it is easier to pass damaged metadata around for inspection.

[Bugzilla:2175198](#)

New per-device counter is added for SCSI devices

A new per-device counter, **iotmo_cnt**, is now added for the I/O timeouts in the SCSI updates. In addition to the **iorequest_cnt** count of I/O requests, the **iodone_cnt** I/O completions, and the **ioerr_cnt** I/O errors, the number of request timeouts can be seen. For example:

```
/sys/devices/pci0000:16/0000:16:02.0/0000:17:00.0/host2/target2:2:0/2:2:0:0/iorequest_cnt
/sys/devices/pci0000:16/0000:16:02.0/0000:17:00.0/host2/target2:2:0/2:2:0:0/iodone_cnt
/sys/devices/pci0000:16/0000:16:02.0/0000:17:00.0/host2/target2:2:0/2:2:0:0/iotmo_cnt
/sys/devices/pci0000:16/0000:16:02.0/0000:17:00.0/host2/target2:2:0/2:2:0:0/ioerr_cnt
```

[Bugzilla:2171093^{\[1\]}](#)

mpathcleanup flushes the multipath devices in device-mapper-multipath

The **mpathcleanup** tool works on SCSI-based multipath devices and removes the multipath device along with the SCSI path devices. Some users need to remove multipath devices and their path devices regularly. Previously, there was no tool available to remove multipath devices and a user-defined script was required for this operation.

With this new tool, users can now easily remove multipath devices and their underlying storage, and there is no need to create any script for this operation.

[Jira:RHEL-782^{\[1\]}](#)

nvme-cli rebased to version 2.4

The **nvme-cli** package has been upgraded to version 2.4, which provides multiple bug fixes and enhancements. Notable changes include:

- Supports TLS over TCP.

- Fixes incorrect ordering of the **systemd** auto-connect services to mount file systems using the **/etc/fstab** file.
- Fixes printing of the **u32** values.
- Validates storage tag size correctly.
- Supports the **nvme effects-log** command for fabrics controllers.

[Bugzilla:2159929^{\[1\]}](#)

4.11. HIGH AVAILABILITY AND CLUSTERS

Support for failover of LVM volume groups with missing physical volumes

The **LVM-activate** resource agent now supports two new options that allow volume group failover if the volume group is missing physical volumes:

- The **majoritypvs** option allows the system ID to be changed on a volume group when a volume group is missing physical volumes, provided that a majority of physical volumes are present.
- The **degraded_activation** option allows RAID logical volumes in a volume group to be activated when legs are missing, provided that sufficient devices are available for RAID to provide all the data in the logical volume.

[Bugzilla:2174911^{\[1\]}](#)

IPaddr2 and IPsrcaddr cluster resource agents now support policy-based routing

The **IPaddr2** and **IPsrcaddr** cluster resource agents now support policy-based routing, which enables you to configure complex routing scenarios. Policy-based routing requires that you configure the resource agent's **table** parameter.

[Bugzilla:2142518](#)

The Filesystem resource agent now supports the EFS file system type

The **ocf:heartbeat:Filesystem** cluster resource agent now supports the Amazon Elastic File System (EFS). You can now specify **fstype=efs** when configuring a **Filesystem** resource.

[Bugzilla:2142002](#)

New pcs parsing requires meta keyword when specifying clone meta attributes

To ensure consistency in the **pcs** command format, configuring clone meta attributes with the **pcs resource clone**, **pcs resource promotable**, and **pcs resource create** commands without specifying the **meta** keyword is now deprecated.

Previously, the **meta** keyword was ignored in the **pcs resource clone** and **pcs resource promotable** commands. In the **pcs resource create** command, however, the meta attributes specified after the **meta** keyword when it followed the **clone** keyword were assigned to the resource rather than to the clone. With this updated parsing algorithm, meta attributes specified after the **meta** keyword when it follows the **clone** keyword are assigned to the clone. To maintain compatibility with existing scripts which rely on the older format, you must specify the **--future** command option to enable this new argument processing when creating a cloned resource with the **pcs resource create** command.

The following command now creates a resource with the meta attribute **mv=v1** and a clone with the meta attribute **mv=v2**:

```
pcs resource create dummy1 ocf:pacemaker:Dummy meta m1=v1 clone meta m2=v2 --future
```

[Bugzilla:2168155](#)

Displaying the **pcs** commands for re-creating configured resource constraints

You can now display the **pcs constraint** commands that can be used to re-create configured resource constraints on a different system by using the **pcs constraint** command with the new **--output-format=cmd** option. The default output format is plain text, as in previous releases, which you can specify with the **--output-format=text** option. The plain text format has been changed slightly to make it consistent with the output format of other **pcs** commands.

[Bugzilla:2163953](#)

Rebase Pacemaker packages to version: 2.1.6

The Pacemaker packages have been upgraded to upstream version 2.1.6, which provides several enhancements and bug fixes over the previous version.

The following features have been added:

- Previously, when a Pacemaker Remote connection was lost, Pacemaker would always purge its transient node attributes. This was unnecessary if the connection was quickly recoverable and the remote daemon had not restarted in the meantime. Pacemaker Remote nodes now preserve transient node attributes after a brief, recoverable connection outage.
- The **alert_snmp.sh.sample** alert agent, which is the sample alert agent provided with Pacemaker, now supports the SNMPv3 protocol and SNMPv2. With this update, you can copy the **alert_snmp.sh.sample** agent without modification to use SNMPv3 with Pacemaker alerts.
- Pacemaker alerts and alert recipients now support an **enabled** meta option. Setting this option to **false** for an alert disables the alert. Setting this option to **true** for an alert and **false** for a particular recipient disables the alert for that recipient. The default value for this option is **true**. You can use this option to temporarily disable an alert for any reason, such as planned maintenance.

The following bugs have been fixed:

- Pacemaker Designated Controller elections no longer finalized until all pending actions are complete and no action results are lost.
- The **fence_scsi** agent is now able to auto-detect shared **lvmlockd** devices when the **devices** attribute is not set.
- Resource stickiness now properly compares against colocation scores.
- The **crm_resource** command now allows banning or moving a bundle with only a single active replica.
- Previously, promotable clone instances were assigned in numerical order, with promoted instances first. As a result, if a promoted clone instance needed to start, an unpromoted instance in some cases restarted unexpectedly, because the instance numbers changed. With this fix, roles are considered when assigning instance numbers to nodes and as a result no unnecessary restarts occur.

[Bugzilla:2189301](#)

Enhancements to the `pcs property` command

The `pcs property` command now supports the following enhancements:

- The `pcs property config --output-format=` option
 - Specify `--output-format=cmd` to display the `pcs property set` command created from the current cluster properties configuration. You can use this command to re-create configured cluster properties on a different system.
 - Specify `--output-format=json` to display the configured cluster properties in JSON format.
 - Specify `output-format=text` to display the configured cluster properties in plain text format, which is the default value for this option.
- The `pcs property defaults` command, which replaces the deprecated `pcs property --defaults` option
- The `pcs property describe` command, which describes the meaning of cluster properties

[Bugzilla:2163914](#)

4.12. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

A new environment variable in Python to control parsing of email addresses

To mitigate [CVE-2023-27043](#), a backward incompatible change to ensure stricter parsing of email addresses was introduced in Python 3.

The update in [RHSA-2024:2024](#) introduces a new `PYTHON_EMAIL_DISABLE_STRICT_ADDR_PARSING` environment variable. When you set this variable to `true`, the previous, less strict parsing behavior is the default for the entire system:

```
export PYTHON_EMAIL_DISABLE_STRICT_ADDR_PARSING=true
```

However, individual calls to the affected functions can still enable stricter behavior.

You can achieve the same result by creating the `/etc/python/email.cfg` configuration file with the following content:

```
[email_addr_parsing]
PYTHON_EMAIL_DISABLE_STRICT_ADDR_PARSING = true
```

For more information, see the Knowledgebase article [Mitigation of CVE-2023-27043 introducing stricter parsing of email addresses in Python](#).

Jira:[RHELDOCS-17369](#)^[1]

A new `nodejs:20` module stream is fully supported

A new module stream, `nodejs:20`, previously available as a Technology Preview, is fully supported with the release of the [RHEA-2023:7252](#) advisory. The `nodejs:20` module stream now provides **Node.js 20.9**, which is a Long Term Support (LTS) version.

Node.js 20 included in RHEL 9.3 provides numerous new features, bug fixes, security fixes, and performance improvements over **Node.js 18** available since RHEL 9.1.

Notable changes include:

- The **V8** JavaScript engine has been upgraded to version 11.3.
- The **npm** package manager has been upgraded to version 9.8.0.
- **Node.js** introduces a new experimental Permission Model.
- **Node.js** introduces a new experimental Single Executable Application (SEA) feature.
- **Node.js** provides improvements to the Experimental ECMAScript modules (ESM) loader.
- The native test runner, introduced as an experimental **node:test** module in **Node.js 18**, is now considered stable.
- **Node.js** provides various performance improvements.

To install the **nodejs:20** module stream, use:

```
# dnf module install nodejs:20
```

If you want to upgrade from the **nodejs:18** stream, see [Switching to a later stream](#).

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

[Bugzilla:2186717](#)

A new filter argument to the Python tarfile extraction functions

To mitigate [CVE-2007-4559](#), Python adds a **filter** argument to the **tarfile** extraction functions. The argument allows turning **tar** features off for increased safety (including blocking the CVE-2007-4559 directory traversal attack). If a filter is not specified, the **'data'** filter, which is the safest but most limited, is used by default in RHEL. In addition, Python emits a warning when your application has been affected.

For more information, including instructions to hide the warning, see the Knowledgebase article [Mitigation of directory traversal attack in the Python tarfile library \(CVE-2007-4559\)](#).

[Jira:RHELDPCS-16405^{\[1\]}](#)

The HTTP::Tiny Perl module now verifies TLS certificates by default

The default value for the **verify_SSL** option in the **HTTP::Tiny** Perl module has been changed from **0** to **1** to verify TLS certificates when using HTTPS. This change fixes [CVE-2023-31486](#) for **HTTP::Tiny** and [CVE-2023-31484](#) for the CPAN Perl module.

To make support for TLS verification available, this update adds the following dependencies to the **perl-HTTP-Tiny** package:

- **perl-IO-Socket-SSL**
- **perl-Mozilla-CA**
- **perl-Net-SSLeay**

Bugzilla:2228412^[1]

httpd rebased to version 2.4.57

The Apache HTTP Server has been updated to version 2.4.57, which provides bug fixes, enhancements, and security fixes over version 2.4.53 available since RHEL 9.1.

Notable enhancements include:

- The **rotatelogs** utility provided with **httpd** introduces a new **-T** option to truncate all rotated logfiles except the initial log file.
- The **LDAPConnectionPoolTTL** directive of the **modldap** module now accepts negative values to enable reuse of connections of any age. Previously, a negative value was handled as an error.
- Workers from the **mod_proxy_hcheck** module now correctly time out according to the worker timeout settings.
- The **hcmethod** parameter of the **mod_proxy_hcheck** module now provides new **GET11**, **HEAD11**, and **OPTIONS11** methods for HTTP/1.1 requests.

Bugzilla:2184403

A new mod_authnz_fcgi module in httpd

The Apache HTTP Server now includes the **mod_authnz_fcgi** module, which enables FastCGI authorizer applications to authenticate users and authorize access to resources.

The **mod_authnz_fcgi** module is not loaded by default. To load this module, uncomment the following line in the **/etc/httpd/conf.modules.d/00-optional.conf** file:

```
LoadModule authnz_fcgi_module modules/mod_authnz_fcgi.so
```

Bugzilla:2173295^[1]

A new ssl_pass_phrase_dialog directive in nginx:1.22

With this update to the **nginx:1.22** module stream, you can use the new **ssl_pass_phrase_dialog** directive to configure an external program that is called at **nginx** start for each encrypted private key.

To use the new directive, add one of the following lines to the **/etc/nginx/nginx.conf** file:

- To call an external program for each encrypted private key file, enter:

```
ssl_pass_phrase_dialog exec:<path_to_program>;
```

nginx calls this program with the following two arguments:

- The server name specified in the **server_name** setting.
 - One of the following algorithms: **RSA**, **DSA**, **EC**, **DH**, or **UNK** if a cryptographic algorithm cannot be recognized.
- If you want to manually enter a passphrase for each encrypted private key file, enter:

```
ssl_pass_phrase_dialog builtin;
```

This is the default behavior if **ssl_pass_phrase_dialog** is not configured.

Note that the **nginx** service fails to start if you use this method but have at least one private key protected by a passphrase. In this case, use one of the other methods.

- If you want **systemd** to prompt for the passphrase for each encrypted private key when you start the **nginx** service by using the **systemctl** utility, enter:

```
ssl_pass_phrase_dialog exec:/usr/libexec/nginx-ssl-pass-dialog;
```

Note that the **ssl_pass_phrase_dialog** directive in **nginx** is similar to the **SSLPassPhraseDialog** directive in the Apache HTTP Server.

[Bugzilla:2170808](#)

A new rhel9/squid container image

The **rhel9/squid** container image is now available in the Red Hat Container Registry. **Squid** is a high-performance proxy caching server for web clients, supporting FTP, gopher, and HTTP data objects. Unlike traditional caching software, **Squid** handles all requests in a single, non-blocking, I/O-driven process. **Squid** keeps metadata and especially hot objects cached in RAM, caches DNS lookups, supports non-blocking DNS lookups, and implements negative caching of failed requests.

To pull the new container image, run:

```
# podman pull registry.redhat.io/rhel9/squid
```

[Bugzilla:2178953](#)

A new module stream: redis:7

Redis 7, an advanced key-value store, is now available as a new module stream, **redis:7**.

Notable changes over **Redis 6** include:

- Server-side scripting in the Redis Functions API
- Fine-grained access control list (ACL) support
- Shared publish/subscribe (pub/sub) support for clusters
- Various new commands and command arguments

Redis 7 introduces several backward incompatible changes, for example:

- **Redis 7** now stores append-only files (AOF) as multiple files in a folder
- **Redis 7** uses a new version format for Redis Database (RDB) files that is incompatible with earlier versions

For a complete list of features and incompatible changes, see the [upstream release notes](#).

To install the **redis:7** module stream, use:

```
# dnf module install redis:7
```


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

[Bugzilla:2129826](#)

4.13. COMPILERS AND DEVELOPMENT TOOLS

A new **glibc** option to influence optimized routine usage on IBM Z

On the IBM Z architecture, the **glibc** library selects function implementations based on the hardware capabilities, such as **hwcaps** and **stfle** bits. With this update, you can direct the choice made by the library by setting the **glibc.cpu.hwcaps** tunable.

[Bugzilla:2169978^{\[1\]}](#)

Improved string and memory routine performance on Intel® Xeon® v5-based hardware in **glibc**

Previously, the default amount of cache used by **glibc** for string and memory routines resulted in lower than expected performance on Intel® Xeon® v5-based systems. With this update, the amount of cache to use has been tuned to improve performance.

[Bugzilla:2213907](#)

The system GCC compiler updated to version 11.4.1

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

The system GCC compiler has been updated to version 11.4.1, which includes numerous bug fixes and enhancements available in the upstream GCC.

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

[Bugzilla:2193180](#)

GCC now supports preserving register arguments

With this update, you can now store argument register content to the stack and generate proper Call Frame Information (CFI) to allow the unwinder to locate it without negatively impacting performance.

[Bugzilla:2168204^{\[1\]}](#)

A new **-mdaz-ftz** option in GCC on the 64-bit Intel architecture

The system version of GNU Compiler Collection (GCC) on the 64-bit Intel architecture now supports the **-mdaz-ftz** option to enable flush-to-zero (FTZ) and denormals-are-zero (DAZ) flags in the MXCSR Control and Status Register.

[Bugzilla:2208908](#)

New GCC Toolset 13

GCC Toolset 13 is a compiler toolset that provides recent versions of development tools. It is available as an Application Stream in the form of a Software Collection in the AppStream repository.

The GCC compiler has been updated to version 13.1.1, which provides many bug fixes and enhancements that are available in upstream GCC.

The following tools and versions are provided by GCC Toolset 13:

Tool	Version
GCC	13.1.1
GDB	12.1
binutils	2.40
dwz	0.14
annobin	12.20

To install GCC Toolset 13, run the following command as root:

```
# dnf install gcc-toolset-13
```

To run a tool from GCC Toolset 13:

```
$ scl enable gcc-toolset-13 tool
```

To run a shell session where tool versions from GCC Toolset 13 override system versions of these tools:

```
$ scl enable gcc-toolset-13 bash
```

For more information, see [GCC Toolset 13](#) and [Using GCC Toolset](#).

Bugzilla:2171919^[1], Bugzilla:2171930

GCC Toolset 13: GCC rebased to version 13.1.1

In GCC Toolset 13, the GNU Compiler Collection (GCC) has been updated to version 13.1.1. Notable changes include:

General improvements

- OpenMP:
 - OpenMP 5.0: Fortran now supports some non-rectangular loop nests. Such support was added for C/C++ in GCC 11.
 - Many OpenMP 5.1 features have been added.
 - Initial support for OpenMP 5.2 features has been added.
- A new debug info compression option value, **-gz=zstd**, is now available.
- The **-Ofast**, **-ffast-math**, and **-funsafe-math-optimizations** options no longer add startup code to alter the floating-point environment when producing a shared object with the **-shared** option.

- GCC can now emit its diagnostics using Static Analysis Results Interchange Format (SARIF), a JSON-based format suited for capturing the results of static analysis tools (such as GCC's **-fanalyzer**). You can also use SARIF to capture other GCC warnings and errors in a machine-readable format.
- Link-time optimization improvements have been implemented.

New languages and language-specific improvements

C family:

- A new **-Wxor-used-as-pow** option warns about uses of the exclusive or (`^`) operator where the user might have meant exponentiation.
- Three new function attributes have been added for documenting **int** arguments that are file descriptors:
 - **`attribute((fd_arg(N)))`**
 - **`attribute((fd_arg_read(N)))`**
 - **`attribute((fd_arg_write(N)))`**

These attributes are also used by **-fanalyzer** to detect misuses of file descriptors.

- A new statement attribute, **`attribute((assume(EXPR)))`**, has been added for C++23 portable assumptions. The attribute is supported also in C or earlier C++.
- GCC can now control when to treat the trailing array of a structure as a flexible array member for the purpose of accessing the elements of such an array. By default, all trailing arrays in aggregates are treated as flexible array members. Use the new command-line option **-fstrict-flex-arrays** to control what array members are treated as flexible arrays.

C:

- Several C23 features have been implemented:
 - Introduced the **`nullptr`** constant.
 - Enumerations enhanced to specify underlying types.
 - Requirements for variadic parameter lists have been relaxed.
 - Introduced the **`auto`** feature to enable type inference for object definitions.
 - Introduced the **`constexpr`** specifier for object definitions.
 - Introduced storage-class specifiers for compound literals.
 - Introduced the **`typeof`** object (previously supported as an extension) and the **`typeof_unqual`** object.
 - Added new keywords: **`alignas`**, **`alignof`**, **`bool`**, **`false`**, **`static_assert`**, **`thread_local`**, and **`true`**.
 - Added the **`[[noreturn]]`** attribute to specify that a function does not return execution to its caller.
 - Added support for empty initializer braces.

- Added support for **STDC_VERSION_*_H** header version macros.
- Removed the **ATOMIC_VAR_INIT** macro.
- Added the **unreachable** macro for the **<stddef.h>** header.
- Removed trigraphs.
- Removed unprototyped functions.
- Added **printf** and **scanf** format checking through the **-Wformat** option for the **%wN** and **%wfN** format length modifiers.
- Added support for identifier syntax of Unicode Standard Annex (UAX) 31.
- Existing features adopted in C23 have been adjusted to follow C23 requirements and are not diagnosed using the **-std=c2x -Wpedantic** option.
- A new **-Wenum-int-mismatch** option warns about mismatches between an enumerated type and an integer type.

C++:

- Implemented excess precision support through the **-fexcess-precision** option. It is enabled by default in strict standard modes such as **-std=c++17**, where it defaults to **-fexcess-precision=standard**. In GNU standard modes such as **-std=gnu++20**, it defaults to **-fexcess-precision=fast**, which restores previous behavior.
The **-fexcess-precision** option affects the following architectures:
 - Intel 32- and 64-bit using x87 math, in some cases on Motorola 68000, where **float** and **double** expressions are evaluated in **long double** precision.
 - 64-bit IBM Z systems where **float** expressions are evaluated in **double** precision.
 - Several architectures that support the **std::float16_t** or **std::bfloat16_t** types, where these types are evaluated in **float** precision.
- Improved experimental support for C++23, including:
 - Added support for labels at the end of compound statements.
 - Added a type trait to detect reference binding to a temporary.
 - Reintroduced support for volatile compound operations.
 - Added support for the **#warning** directive.
 - Added support for delimited escape sequences.
 - Added support for named universal character escapes.
 - Added a compatibility and portability fix for the **char8_t** type.
 - Added static **operator()** function objects.
 - Simplified implicit moves.
 - Rewriting equality in expressions is now less of a breaking change.

- Removed non-encodable wide character literals and wide multicharacter literals.
- Relaxed some **constexpr** function restrictions.
- Extended floating-point types and standard names.
- Implemented portable assumptions.
- Added support for UTF-8 as a portable source file encoding standard.
- Added support for static **operator[]** subscripts.
- New warnings:
 - **-Wself-move** warns when a value is moved to itself with **std::move**.
 - **-Wdangling-reference** warns when a reference is bound to a temporary whose lifetime has ended.
 - The **-Wpessimizing-move** and **-Wredundant-move** warnings have been extended to warn in more contexts.
- The new **-nostdlib++** option enables linking with **g++** without implicitly linking in the C++ standard library.

Changes in the **libstdc++** runtime library

- Improved experimental support for C++20, including:
 - Added the **<format>** header and the **std::format** function.
 - Added support in the **<chrono>** header for the **std::chrono::utc_clock** clock, other clocks, time zones, and the **std::format** function.
- Improved experimental support for C++23, including:
 - Additions to the **<ranges>** header: **views::zip**, **views::zip_transform**, **views::adjacent**, **views::adjacent_transform**, **views::pairwise**, **views::slide**, **views::chunk**, **views::chunk_by**, **views::repeat**, **views::chunk_by**, **views::cartesian_product**, **views::as_rvalue**, **views::enumerate**, **views::as_const**.
 - Additions to the **<algorithm>** header: **ranges::contains**, **ranges::contains_subrange**, **ranges::iota**, **ranges::find_last**, **ranges::find_last_if**, **ranges::find_last_if_not**, **ranges::fold_left**, **ranges::fold_left_first**, **ranges::fold_right**, **ranges::fold_right_last**, **ranges::fold_left_with_iter**, **ranges::fold_left_first_with_iter**.
 - Support for monadic operations for the **std::expected** class template.
 - Added **constexpr** modifiers to the **std::bitset**, **std::to_chars** and **std::from_chars** functions.
 - Added library support for extended floating-point types.
- Added support for the **<experimental/scope>** header from version 3 of the Library Fundamentals Technical Specification (TS).
- Added support for the **<experimental/synchronized_value>** header from version 2 of the Concurrency TS.

- Added support for many previously unavailable features in freestanding mode. For example:
 - The **std::tuple** class template is now available for freestanding compilation.
 - The **libstdc++** library adds components to the freestanding subset, such as **std::array** and **std::string_view**.
 - The **libstdc++** library now respects the **-ffreestanding** compiler option, so it is no longer necessary to build a separate freestanding installation of the **libstdc++** library. Compiling with **-ffreestanding** will restrict the available features to the freestanding subset, even if the **libstdc++** library was built as a full, hosted implementation.

New targets and target-specific Improvements

The 64-bit ARM architecture:

- Added support for the **armv9.1-a**, **armv9.2-a**, and **armv9.3-a** arguments for the **-march=** option.

The 32- and 64-bit AMD and Intel architectures:

- For both C and C++, the **__bf16** type is supported on systems with Streaming SIMD Extensions 2 and above enabled.
- The real **__bf16** type is now used for **AVX512BF16** instruction intrinsics. Previously, **__bfloat16**, a typedef of short, was used. Adjust your **AVX512BF16** related source code when upgrading GCC 12 to GCC 13.
- Added new Instruction Set Architecture (ISA) extensions to support the following Intel instructions:
 - **AVX-IFMA** whose instruction intrinsics are available through the **-mavxifma** compiler switch.
 - **AVX-VNNI-INT8** whose instruction intrinsics are available through the **-mavxvnniint8** compiler switch.
 - **AVX-NE-CONVERT** whose instruction intrinsics are available through the **-mavxneconvert** compiler switch.
 - **CMPccXADD** whose instruction intrinsics are available through the **-mcmpccxadd** compiler switch.
 - **AMX-FP16** whose instruction intrinsics are available through the **-mamx-fp16** compiler switch.
 - **PREFETCHI** whose instruction intrinsics are available through the **-mprefetchi** compiler switch.
 - **RAO-INT** whose instruction intrinsics are available through the **-mraoint** compiler switch.
 - **AMX-COMPLEX** whose instruction intrinsics are available through the **-mamx-complex** compiler switch.
- GCC now supports AMD CPUs based on the **znver4** core through the **-march=znver4** compiler switch. The switch makes GCC consider using 512-bit vectors when auto-vectorizing.

Improvements to the static analyzer

- The static analyzer has gained 20 new warnings:
 - **-Wanalyzer-allocation-size**
 - **-Wanalyzer-deref-before-check**
 - **-Wanalyzer-exposure-through-uninit-copy**
 - **-Wanalyzer-imprecise-fp-arithmetic**
 - **-Wanalyzer-infinite-recursion**
 - **-Wanalyzer-jump-through-null**
 - **-Wanalyzer-out-of-bounds**
 - **-Wanalyzer-putenv-of-auto-var**
 - **-Wanalyzer-tainted-assertion**
 - Seven new warnings relating to misuse of file descriptors:
 - **-Wanalyzer-fd-access-mode-mismatch**
 - **-Wanalyzer-fd-double-close**
 - **-Wanalyzer-fd-leak**
 - **-Wanalyzer-fd-phase-mismatch** (for example, calling **accept** on a socket before calling **listen** on it)
 - **-Wanalyzer-fd-type-mismatch** (for example, using a stream socket operation on a datagram socket)
 - **-Wanalyzer-fd-use-after-close**
 - **-Wanalyzer-fd-use-without-check**
 - Also implemented special-casing handling of the behavior of the **open**, **close**, **creat**, **dup**, **dup2**, **dup3**, **pipe**, **pipe2**, **read**, and **write** functions.
 - Four new warnings for misuses of the `<stdarg.h>` header:
 - **-Wanalyzer-va-list-leak** warns about missing a **va_end** macro after a **va_start** or **va_copy** macro.
 - **-Wanalyzer-va-list-use-after-va-end** warns about a **va_arg** or **va_copy** macro used on a **va_list** object type that has had the **va_end** macro called on it.
 - **-Wanalyzer-va-arg-type-mismatch** type-checks **va_arg** macro usage in interprocedural execution paths against the types of the parameters that were actually passed to the variadic call.
 - **-Wanalyzer-va-list-exhausted** warns if a **va_arg** macro is used too many times on a **va_list** object type in interprocedural execution paths.
- Numerous other improvements.

Backwards incompatible changes

For C++, construction of global iostream objects such as **std::cout**, **std::cin** is now done inside the standard library, instead of in every source file that includes the `<iostream>` header. This change improves the startup performance of C++ programs, but it means that code compiled with GCC 13.1 will crash if the correct version of **libstdc++.so** is not used at runtime. See the [documentation](#) about using the correct **libstdc++.so** at runtime. Future GCC releases will mitigate the problem so that the program cannot be run at all with an earlier incompatible **libstdc++.so**.

Bugzilla:2172093^[1]

GCC Toolset 13: annobin rebased to version 12.20

GCC Toolset 13 provides the **annobin** package version 12.20. Notable enhancements include:

- Added support for moving **annobin** notes into a separate debug info file. This results in reduced executable binary size.
- Added support for a new smaller note format reduces the size of the separate debuginfo files and the time taken to create these files.

Bugzilla:2171923^[1]

GCC Toolset 13: GDB rebased to version 12.1

GCC Toolset 13 provides GDB version 12.1.

Notable bug fixes and enhancements include:

- GDB now styles source code and disassembler by default. If styling interferes with automation or scripting of GDB, you can disable it by using the **maint set gnu-source-highlight enabled off** and **maint set style disassembler enabled off** commands.
- GDB now displays backtraces whenever it encounters an internal error. If this affects scripts or automation, you can use the **maint set backtrace-on-fatal-signal off** command to disable this feature.

C/C++ improvements:

- GDB now treats functions or types involving C++ templates similarly to function overloads. You can omit parameter lists to set breakpoints on families of template functions, including types or functions composed of multiple template types. **Tab** completion has gained similar improvements.

Terminal user interface (TUI):

- **tui layout**
tui focus

tui refresh

tui window height

These are the new names for the old **layout**, **focus**, **refresh**, and **winheight** TUI commands. The old names still exist as aliases to these new commands.

- **tui window width**
winwidth

Use the new **tui window width** command, or the **winwidth** alias, to adjust the width of a TUI window when windows are laid out in horizontal mode.

- **info win**

This command now includes information about the width of the TUI windows in its output.

Machine Interface (MI) changes:

- The default version of the MI interpreter is now 4 (`-i=mi4`).
- The **-add-inferior** command with no flag now inherits the connection of the current inferior. This restores the behavior of GDB before version 10.
- The **-add-inferior** command now accepts a **--no-connection** flag that causes the new inferior to start without a connection.
- The **script** field in breakpoint output (which is syntactically incorrect in MI 3 and earlier) has become a list in MI 4. This affects the following commands and events:
 - **-break-insert**
 - **-break-info**
 - **=breakpoint-created**
 - **=breakpoint-modified**
Use the **-fix-breakpoint-script-output** command to enable the new behavior with earlier MI versions.

New commands:

- **maint set internal-error backtrace [on|off]**
maint show internal-error backtrace

maint set internal-warning backtrace [on|off]

maint show internal-warning backtrace

GDB can now print a backtrace of itself when it encounters internal error or internal warning. This is enabled by default for internal errors and disabled by default for internal warnings.

- **exit**

You can exit GDB using the new **exit** command in addition to the existing **quit** command.

- **maint set gnu-source-highlight enabled [on|off]**
maint show gnu-source-highlight enabled

Enables or disables the GNU Source Highlight library for adding styling to source code. When disabled, the library is not used even if it is available. When the GNU Source Highlight library is not used the Python Pygments library is used instead.

- **set suppress-cli-notifications [on|off]**
show suppress-cli-notifications

Controls if printing the notifications is suppressed for CLI or not. CLI notifications occur when you change the selected context (such as the current inferior, thread, or frame), or when the program being debugged stops (for example: because of hitting a breakpoint, completing source-stepping, or an interrupt).

- **set style disassembler enabled [on|off]**
show style disassembler enabled

When enabled, the command applies styling to disassembler output if GDB is compiled with Python support and the Python Pygments package is available.

Changed commands:

- **set logging [on|off]**
Deprecated and replaced by the **set logging enabled [on|off]** command.
- **print**
Printing of floating-point values with base-modifying formats such as `/x` has been changed to display the underlying bytes of the value in the required base.
- **clone-inferior**
The **clone-inferior** command now ensures that the **TTY**, **CMD**, and **ARGs** settings are copied from the original inferior to the new one. All modifications to the environment variables done using the **set environment** or **unset environment** commands are also copied to the new inferior.

Python API:

- The new **`gdb.add_history()`** function takes a **`gdb.Value`** object and adds the value it represents to GDB's history list. The function returns an integer, which is the index of the new item in the history list.
- The new **`gdb.history_count()`** function returns the number of values in GDB's value history.
- The new **`gdb.events.gdb_exiting`** event is called with a **`gdb.GdbExitingEvent`** object that has the read-only attribute **`exit_code`** containing the value of the GDB exit code. This event is triggered before GDB's exit before GDB starts to clean up its internal state.
- The new **`gdb.architecture_names()`** function returns a list containing all of the possible **`Architecture.name()`** values. Each entry is a string.
- The new **`gdb.Architecture.integer_type()`** function returns an integer type given a size and a signed-ness.
- The new **`gdb.TargetConnection`** object type represents a connection (as displayed by the **`info connections`** command). A sub-class, **`gdb.RemoteTargetConnection`**, represents **`remote`** and **`extended-remote`** connections.
- The **`gdb.Inferior`** type now has a **`connection`** property that is an instance of the **`gdb.TargetConnection`** object, the connection used by this inferior. This can be **`None`** if the inferior has no connection.
- The new **`gdb.events.connection_removed`** event registry emits a **`gdb.ConnectionEvent`** event when a connection is removed from GDB. This event has a **`connection`** property, a **`gdb.TargetConnection`** object for the connection being removed.
- The new **`gdb.connections()`** function returns a list of all currently active connections.
- The new **`gdb.RemoteTargetConnection.send_packet(PACKET)`** method is equivalent to the existing **`maint packet`** CLI command. You can use it to send a specified packet to the remote target.

- The new `gdb.host_charset()` function returns the name of the current host character set as a string.
- The new `gdb.set_parameter(NAME, VALUE)` function sets the GDB parameter `NAME` to `VALUE`.
- The new `gdb.with_parameter(NAME, VALUE)` function returns a context manager that temporarily sets the GDB parameter `NAME` to `VALUE` and then resets it when the context is exited.
- The `gdb.Value.format_string` method now takes a `styling` argument, which is a boolean. When `true`, the returned string can include escape sequences to apply styling. The styling is present only if styling is turned on in GDB (see [help set styling](#)). When `false`, which is the default if the `styling` argument is not given, no styling is applied to the returned string.
- The new read-only attribute `gdb.InferiorThread.details` is either a string containing additional target-specific thread-state information, or `None` if there is no such additional information.
- The new read-only attribute `gdb.Type.is_scalar` is `True` for scalar types, and `False` for all other types.
- The new read-only attribute `gdb.Type.is_signed` should only be read when `Type.is_scalar` is `True`, and will be `True` for signed types and `False` for all other types. Attempting to read this attribute for non-scalar types will raise a `ValueError`.
- You can now add GDB and MI commands implemented in Python.

For more information see the upstream release notes:

[What has changed in GDB?](#)

Bugzilla:2172096^[1]

GCC Toolset 13: bintuils rebased to version 2.40

GCC Toolset 13 provides the `binutils` package version 2.40. Notable enhancements include:

Linkers:

- The new `-w` (`--no-warnings`) command-line option for the linker suppresses the generation of any warning or error messages. This is useful in case you need to create a known non-working binary.
- The ELF linker now generates a warning message if:
 - The stack is made executable
 - It creates a memory resident segment with all three of the `Read`, `Write` and `eXecute` permissions set
 - It creates a thread local data segment with the `eXecute` permission set. You can disable these warnings by using the `--no-warn-exec-stack` or `--no-warn-rwx-segments` options.
- The linker can now insert arbitrary JSON-format metadata into binaries that it creates.

Other tools:

- A new the **objdump** tool's **--private** option to display fields in the file header and section headers for Portable Executable (PE) format files.
- A new **--strip-section-headers** command-line option for the **objcopy** and **strip** utilities to remove the ELF section header from ELF files.
- A new **--show-all-symbols** command-line option for the **objdump** utility to display all symbols that match a given address when disassembling, as opposed to the default function of displaying only the first symbol that matches an address.
- A new **-W** (**--no-weak**) option to the **nm** utility to make it ignore weak symbols.
- The **objdump** utility now supports syntax highlighting of disassembler output for some architectures. Use the **--disassembler-color=MODE** command-line option, with *MODE* being one of the following:
 - **off**
 - **color** - This option is supported by all terminal emulators.
 - **extended-color** - This option uses 8-bit colors not supported by all terminal emulators.

[Bugzilla:2171926^{\[1\]}](#)

libabigail rebased to version 2.3

The **libabigail** package has been updated to version 2.3. Notable improvements include:

- The BTF debuginfo format is now supported.
- Improved support for Ada range types.
- A new **[allow_type]** directive in suppression specifications is now supported.
- Added various new properties for the **[supress_type]** suppression specification.
- The ABIXML file format has been updated to version 2.2.
- The SONAME of the library has been changed to reflect its own ABI change.

The **libabigail** package is available in the CodeReady Linux Builder (CRB) repository. Note that packages included in the CodeReady Linux Builder repository are unsupported.

[Bugzilla:2186931](#)

The find-debuginfo script in debugedit now supports the -q (--quiet) flag

With this update, you can use the **find-debuginfo** script's **-q** (**--quiet**) flag in the **debugedit** utility to silence non-error output from the script.

[Bugzilla:2177302](#)

Valgrind rebased to version 3.21.0

Valgrind has been updated to version 3.21.0. Notable enhancements include:

- A new **abexit** value for the **--vgdb-stop-at=event1,event2,...** option notifies the **gdbserver** utility when your program exits abnormally, such as with a nonzero exit code.

- A new **--enable-debuginfod=[yes|no]** option instructs Valgrind to use the **debuginfod** servers listed in the **DEBUGINFOD_URLS** environment variable to fetch any missing DWARF debuginfo information for the program running under Valgrind. The default value for this option is **yes**.



NOTE

The **DEBUGINFOD_URLS** environment variable is not set by default.

- Valgrind now provides GDB Python commands. These GDB front end commands provide a better integration in the GDB command-line interface. Benefits of this are, for example, GDB auto-completion, and command-specific help, searching for a command or command help that matches a regular expression. For relevant monitoring commands, GDB evaluates arguments to simplify usage of monitor commands.
- The **vgdb** utility now supports the extended remote protocol when invoked with the **--multi** option. The GDB **run** command is supported in this mode and, as a result, you can run GDB and Valgrind from a single terminal.
- You can use the **--realloc-zero-bytes-frees=[yes|no]** option to change the behavior of the **realloc()** function with a size of zero for tools that intercept the **malloc()** call.
- The **memcheck** tool now performs checks for the use of the **realloc()** function with a size of zero. Use the new **--show-realloc-size-zero=[yes|no]** switch to disable this feature.
- You can use the new **--history-backtrace-size=value** option for the **helgrind** tool to configure the number of entries to record in the stack traces of earlier accesses.
- The **--cache-sim=[yes|no] cachegrind** option now defaults to **no** and, as a result, only instruction cache read events are gathered by default.
- The source code for the **cg_annotate**, **cg_diff**, and **cg_merge cachegrind** utilities has been rewritten and, as a result, the utilities have more flexible command line option handling. For example, they now support the **--show-percs** and **--no-show-percs** options and the existing **--show-percs=yes** and **--show-percs=no** options.
- The **cg_annotate cachegrind** utility now supports diffing (using the **--diff**, **--mod-filename**, and **--mod-funcname** options) and merging (by passing multiple data files). In addition, **cg_annotate** now provides more information at the file and function level.
- A new user-request for the **DHAT** tool allows you to override the 1024 byte limit on access count histograms for blocks of memory.

The following new architecture-specific instruction sets are now supported:

- 64-bit ARM:
 - v8.2 scalar and vector Floating-point Absolute Difference (FABD), Floating-point Absolute Compare Greater than or Equal (FACGE), Floating-point Absolute Compare Greater Than (FACGT), and Floating-point Add (FADD) instructions.
 - v8.2 Floating-point (FP) compare and conditional compare instructions.
 - Zero variants of v8.2 Floating-point (FP) compare instructions.
- 64-bit IBM Z:

- Support for the **miscellaneous-instruction-extensions facility 3** and the **vector-enhancements facility 2**. This enables programs compiled with the **-march=arch13** or **-march=z15** options to be executed under Valgrind.
- IBM Power:
 - ISA 3.1 support is now complete.
 - ISA 3.0 now supports the deliver a random number (darn) instruction.
 - ISA 3.0 now supports the System Call Vectored (scv) instruction.
 - ISA 3.0 now supports the copy, paste, and cpabort instructions.

[Bugzilla:2124346](#)

systemtap rebased to version 4.9

The **systemtap** package has been upgraded to version 4.9. Notable changes include:

- A new Language-Server-Protocol (LSP) backend for easier interactive drafting of **systemtap** scripts on LSP-capable editors.
- Access to a Python/Jupyter interactive notebook front end.
- Improved handling of DWARF 5 bit fields.

[Bugzilla:2186934](#)

elfutils rebased to version 0.189

The **elfutils** package has been updated to version 0.189. Notable improvements and bug fixes include:

libelf

The **elf_compress** tool now supports the **ELFCOMPRESS_ZSTD** ELF compression type.

libdwfl

The **dwfl_module_return_value_location** function now returns 0 (no return type) for DWARF Information Entries (DIEs) that point to a **DW_TAG_unspecified_type** type tag.

eu-elfcompress

The **-t** and **--type=** options now support the Zstandard (**zstd**) compression format via the **zstd** argument.

[Bugzilla:2182061](#)

libpfm rebased to version 4.13

The **libpfm** package has been updated to version 4.13. With this update, **libpfm** can access performance monitoring hardware native events for the following processor microarchitectures:

- AMD Zen 2
- AMD Zen 3
- AMD Zen 4
- ARM Neoverse N1

- ARM Neoverse N2
- ARM Neoverse V1
- ARM Neoverse V2
- IBM z16
- 4th Generation Intel® Xeon® Scalable Processors

[Bugzilla:2185652](#), [Bugzilla:2047720](#), [Bugzilla:2111940](#), [Bugzilla:2111924](#), [Bugzilla:2111930](#), [Bugzilla:2111933](#), [Bugzilla:2111957](#), [Bugzilla:2111946](#)

papi supports new processor microarchitectures

With this enhancement, you can access performance monitoring hardware using **papi** events presets on the following processor microarchitectures:

- AMD Zen 2
- AMD Zen 3
- ARM Neoverse N1
- ARM Neoverse N2
- ARM Neoverse V1
- ARM Neoverse V2

[Bugzilla:2111923^{\[1\]}](#), [Bugzilla:2111947](#), [Bugzilla:2111942](#)

papi now supports fast performance event count read operations for 64-bit ARM processors

Previously on 64-bit ARM processors, all performance event counter read operations required the use of a resource-intensive system call. **papi** has been updated for 64-bit ARM to let processes monitoring themselves with the performance counters use a faster user-space read of the performance event counters. Setting the `/proc/sys/kernel/perf_user_access` parameter to 1 reduces the average number of clock cycles for **papi** to read 2 counters from 724 cycles to 29 cycles.

[Bugzilla:2186927^{\[1\]}](#)

LLVM Toolset rebased to version 16.0.6

LLVM Toolset has been updated to version 16.0.6.

Notable enhancements include:

- Improvements to optimization
- Support for new CPU extensions
- Improved support for new C++ versions.

Notable backwards incompatible changes include:

- Clang's default C++ standard is now **gnu++17** instead of **gnu++14**.

- The **-Wimplicit-function-declaration**, **-Wimplicit-int** and **-Wincompatible-function-pointer-types** options now default to error for C code. This might affect the behavior of configure scripts.

By default, Clang 16 uses the **libstdc++** library version 13 and **binutils 2.40** provided by GCC Toolset 13.

For more information, see the [LLVM release notes](#) and [Clang release notes](#).

[Bugzilla:2178796](#)

Rust Toolset rebased to version 1.71.1

Rust Toolset has been updated to version 1.71.1. Notable changes include:

- A new implementation of multiple producer, single consumer (mpsc) channels to improve performance
- A new Cargo **sparse** index protocol for more efficient use of the **crates.io** registry
- New **OnceCell** and **OnceLock** types for one-time value initialization
- A new **C-unwind** ABI string to enable usage of forced unwinding across Foreign Function Interface (FFI) boundaries

For more details, see the series of upstream release announcements:

- [Announcing Rust 1.67.0](#)
- [Announcing Rust 1.68.0](#)
- [Announcing Rust 1.69.0](#)
- [Announcing Rust 1.70.0](#)
- [Announcing Rust 1.71.0](#)

[Bugzilla:2191743](#)

The Rust `profiler_builtins` runtime component is now available

With this enhancement, the Rust **profile_builtins** runtime component is now available. This runtime component enables the following compiler options:

-C instrument-coverage

Enables coverage profiling

-C profile-generate

Enables profile-guided optimization

[Bugzilla:2227082^{\[1\]}](#)

Go Toolset rebased to version 1.20.10

Go Toolset has been updated to version 1.20.10.

Notable enhancements include:

- New functions added in the **unsafe** package to handle slices and strings without depending on the internal representation.
- Comparable types can now satisfy comparable constraints.
- A new **crypto/ecdh** package.
- The **go build** and **go test** commands no longer accept the **-i** flag.
- The **go generate** and **go test** commands now accept the **-skip pattern** option.
- The **go build**, **go install**, and other build-related commands now support the **-pgo** and **-cover** flags.
- The **go** command now disables **cgo** by default on systems without a C toolchain.
- The **go version -m** command now supports reading more Go binaries types.
- The **go** command now disables **cgo** by default on systems without a C toolchain.
- Added support for collecting code coverage profiles from applications and integration tests instead of collecting them only from unit tests.

Bugzilla:2185259^[1]

pcp rebased to version 6.0.5

The **pcp** package has been updated to version 6.0.5. Notable changes include:

Collector tool features

- **pmdaproc:**
 - Added support for per-cgroup IRQ PSI metrics in recent kernels
 - Added a new **proc.smaps.pss_dirty** metric
- **pmdasmart:** Added NVME disk information and power state metrics
- **pmdalinux:**
 - Added support for system wide IRQ PSI metrics in recent kernels
 - Added NUMA external memory fragmentation metric
 - Added new networking (TCP, ICMP) metrics
- **pmdaoverhead:** A new PMDA to measure overhead for groups of processes
- **pmdahacluster:** Updated to handle Pacemaker 2.1.5 **crm_mon** output changes

Monitoring tool features

- **pmieconf:**
 - Added support for webhook actions (Event Driven Ansible)
 - Added a new **pmie** rule that checks file descriptor limits

- **pcp2json**: Extended **pcp2json** with an option to send HTTP POSTs
- **pcp-atop**: Added **cgroup**, NUMA memory, and NUMA CPU support
- **pcp-htop**: Added support for a new open file descriptors Meter
- **pcp-ps**: Added capability to show multiple archive samples

[Bugzilla:2175602](#)

PCP's **pmie** utility now supports generating webhook events

The Performance Metrics Inference Engine (**pmie**) utility from Performance Co-Pilot (PCP) now supports generating webhook events. With this update, configured **pmie** rules generate events in a format consumable by Event-Driven Ansible (EDA). As a result, EDA can respond to PCP rules.

To enable this feature, configure all local **pmie** rules to send to a webhook at a given endpoint (URL):

```
# pmieconf modify global webhook_endpoint https://localhost:443/<endpoint>
# pmieconf modify global webhook_action yes
```

[Bugzilla:2185803](#)

grafana rebased to version 9.2.10

The **grafana** package has been updated to version 9.2.10. Notable changes include:

- The heatmap panel is now used throughout Grafana.
- Geomaps can now measure both distance and area.
- The Alertmanager is now based on **Prometheus Alertmanager** version 0.24.
- Grafana Alerting rules now return an **Error** state by default on execution error or timeout.
- Expressions can now be used on public dashboards.
- The join transformation now supports inner joins.
- Public dashboards now allow sharing Grafana dashboards.
- A new Prometheus streaming parser is now available as an opt-in feature.

For more information, see the upstream release notes:

- [What's new in Grafana v9.1](#)
- [What's new in Grafana v9.2](#)

[Bugzilla:2193018](#)

Grafana no longer enables weak cryptographic ciphers

With this update, Grafana no longer enables ciphers that are considered weak for encrypting secure communication. The affected ciphers are:

- **AES128-GCM-SHA256**

- **AES128-SHA**
- **AECDHE-RSA-AES128-SHA**
- **AES256-GCM-SHA384**
- **AES256-SHA**
- **ECDHE-RSA-AES256-SHA**

Bugzilla:2190025^[1]

.NET 8.0 is available

Red Hat Enterprise Linux 9.3 is distributed with .NET version 8.0. Notable improvements include:

- Added support for the C#12 and F#8 language versions.
- Added support for building container images using the .NET Software Development Kit directly.
- Many performance improvements to the garbage collector (GC), Just-In-Time (JIT) compiler, and the base libraries.

Jira:RHELPLAN-164399^[1]

4.14. IDENTITY MANAGEMENT

samba rebased to version 4.18.6

The **samba** packages have been upgraded to upstream version 4.18.6, 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 metadata workloads. This update improves the performance in this scenario.
- The new **wbinfo --change-secret-at=<domain_controller>** command enforces the change of the trust account password on the specified domain controller.
- By default, Samba stores access control lists (ACLs) in the **security.NTACL** extended attribute of files. You can now customize the attribute name with the **acl_xattr:<security_acl_name>** setting in the **/etc/samba/smb.conf** file. Note that a custom extended attribute name is not a protected location as **security.NTACL**. Consequently, users with local access to the server can be able to modify the custom attribute's content and compromise the ACL.

Note that the server message block version 1 (SMB1) protocol has been 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.

Bugzilla:2190415

The **ipaclient** role now allows configuring user subID ranges on the IdM level

With this update, the **ipaclient ansible-freeipa** role provides the **ipaclient_subid** option, using which you can configure subID ranges on the Identity Management (IdM) level. Without the new option set explicitly to **true**, the **ipaclient** role keeps the default behavior and installs the client without subID ranges configured for IdM users.

Previously, the role configured the **sssd authselect** profile that in turn customized the **/etc/nsswitch.conf** file. The subID database did not use IdM and relied only on the local files of **/etc/subuid** and **/etc/subgid**.

[Bugzilla:2175767](#)

Multiple IdM groups and services can now be managed in a single Ansible task

With this enhancement in **ansible-freeipa**, you can add, modify, and delete multiple Identity Management (IdM) user groups and services by using a single Ansible task. For that, use the **groups** and **services** options of the **ipagroup** and **ipaservice** modules.

Using the **groups** option available in **ipagroup**, you can specify multiple group variables that only apply to a particular group. This group is defined by the **name** variable, which is the only mandatory variable for the **groups** option.

Similarly, using the **services** option available in **ipaservice**, you can specify multiple service variables that only apply to a particular service. This service is defined by the **name** variable, which is the only mandatory variable for the **services** option.

[Jira:RHELDOCS-16474^{\[1\]}](#)

ansible-freeipa ipaserver role now supports Random Serial Numbers

With this update, you can use the **ipaserver_random_serial_numbers=true** option with the **ansible-freeipa ipaserver** role. This way, you can generate fully random serial numbers for certificates and requests in PKI when installing an Identity Management (IdM) server using Ansible. With RSNv3, you can avoid range management in large IdM installations and prevent common collisions when reinstalling IdM.



IMPORTANT

RSNv3 is supported only for new IdM installations. If enabled, it is required to use RSNv3 on all PKI services.

[Jira:RHELDOCS-16462^{\[1\]}](#)

ipa rebased to version 4.10.2

The **ipa** package has been upgraded to version 4.10.2. Notable changes include:

- Searching and listing certificates in the IdM CLI and Web UI now offer better performance.

For more information, see the [upstream FreeIPA release notes](#).

[Bugzilla:2196426](#)

The ipaserver_remove_on_server and ipaserver_ignore_topology_disconnect options are now available in the ipaserver role

If removing a replica from an Identity Management (IdM) topology by using the **remove_server_from_domain** option of the **ipaserver ansible-freeipa** role leads to a disconnected topology, you must now specify which part of the domain you want to preserve. Specifically, you must do

the following:

- Specify the **ipaserver_remove_on_server** value to identify which part of the topology you want to preserve.
- Set **ipaserver_ignore_topology_disconnect** to True.

Note that if removing a replica from IdM by using the **remove_server_from_domain** option preserves a connected topology, neither of these options is required.

[Bugzilla:2127903](#)

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:2181465](#)

You can now manage IdM certificates using the **ipacert** Ansible module

You can now use the **ansible-freeipa ipacert** module to request or retrieve SSL certificates for Identity Management (IdM) users, hosts and services. The users, hosts and services can then use these certificates to authenticate to IdM. You can also revoke the certificates, and restore certificates that have been put on hold.

[Bugzilla:2127907](#)

The **optional_pac_tkt_chksum** option helps preserve interoperability between different versions of **krb5**

You can now use the **optional_pac_tkt_chksum** option to preserve the interoperability between RHEL Kerberos Distribution Center (KDC) servers running different versions of the **krb5** package. Specifically, you can change their behavior regarding Privilege Attribute Certificate (PAC) ticket signature verification. If you set the **optional_pac_tkt_chksum** string attribute to **true** for the Kerberos principal expected to sign a ticket, then the KDC does not reject service for user (S4U) requests containing a ticket that lacks the PAC ticket signature. The principal to sign the ticket is the ticket-granting service (TGS) one or a cross-realm TGS one, depending on the realm of the ticket's target service.

Since the **krb5-1.20** release, MIT Kerberos KDCs have required the presence of ticket signatures in PACs based on the encrypted part of Kerberos tickets so that they could process S4U requests successfully. Previously, this was a problem in gradual upgrade scenarios where certain KDCs used **krb5-1.19** or older, while others used **krb5-1.20** or newer. KDCs using the newer versions of **krb5** for S4U requests rejected service tickets that were provided by KDCs using the older versions of **krb5** if a service used them for S4U requests.


For more information about how this feature is used in Identity Management (IdM), see [this pull request](#).

[Bugzilla:2178298](#)

IdM now supports resource-based constrained delegation

With this update, IdM now supports resource-based constrained delegation (RBCD). RBCD allows a granular control of delegation on a resource level and access can be set by the owner of the service to which credentials are delegated.

RBCD can be useful, for example, in an integration between IdM and Active Directory (AD), because AD enforces the use of RBCD when both target and proxy services belong to different forests.



IMPORTANT

Currently, only services in the IdM domain can be configured with RBCD rules. If the target service is part of an AD domain, the permission can only be granted on the AD side. As AD domain controllers cannot resolve IdM service information to create the rule, this is not currently supported.

For more information on delegation scenarios, see the [FreelPA design page](#).

[Bugzilla:2165880](#)

RHEL 9.3 provides 389-ds-base 2.3.4

RHEL 9.3 is distributed with the **389-ds-base** package version 2.3.4. Notable bug fixes and enhancements over version 2.3.4 include:

- <https://www.port389.org/docs/389ds/releases/release-2-2-8.html>
- <https://www.port389.org/docs/389ds/releases/release-2-2-9.html>
- <https://www.port389.org/docs/389ds/releases/release-2-3-0.html>
- <https://www.port389.org/docs/389ds/releases/release-2-3-1.html>
- <https://www.port389.org/docs/389ds/releases/release-2-3-2.html>
- <https://www.port389.org/docs/389ds/releases/release-2-3-3.html>
- <https://www.port389.org/docs/389ds/releases/release-2-3-4.html>

[Bugzilla:2188627](#)

Directory Server can now close a client connection if a **bind** operation fails

Previously, when a **bind** operation failed, some applications that ignore the **bind** return code could load Director Server with further requests.

With the new **nsslapd-close-on-failed-bind** configuration attribute under the **cn=config** entry, the server can close a client connection when the **bind** operation fails. As a result, the server load can be reduced.

[Bugzilla:1987471](#)

Automembership plug-in improvements. It no longer cleans up groups by default

Previously, the automember rebuild task went through all the automember rules and removed all the memberships, then the task rebuilt the memberships from scratch. Thus, the rebuild task was expensive, especially if other **be_txn** plugins are enabled.

With this update, the Automembership plug-in has the following improvements:

- Only one rebuilt task is allowed at a time.
- The Automembership plug-in now does not clean up previous members by default. Use the new CLI option **--cleanup** to intentionally clean up memberships before rebuilding from scratch:

```
# dsconf slapd-instance_name plugins automember fixup -f objectclass=posixaccount -s sub
--cleanup "ou=people,dc=example,dc=com"
```

- Improved logging to show fixup progress.

[Bugzilla:2149025](#)

New **passwordAdminSkipInfoUpdate: on/off** configuration option is now available

You can add a new **passwordAdminSkipInfoUpdate: on/off** setting under the **cn=config** entry to provide a fine grained control over password updates performed by password administrators. When you enable this setting, password updates do not update certain attributes, for example, **passwordHistory**, **passwordExpirationTime**, **passwordRetryCount**, **pwdReset**, and **passwordExpWarned**.

[Bugzilla:2166332](#)

New **slapi_memberof()** plug-in function is now available for Directory Server plug-ins and client applications

The new **slapi_memberof()** function retrieves distinguished names (DNs) of groups to which the given entry belongs directly or indirectly. Previously, MemberOf, Referential Integrity, and ACL plug-ins implemented their own mechanism to retrieve such groups. With this update, you can use the **slapi_memberof()** function that introduces a unified mechanism to return group DN's.

[Bugzilla:2189946](#)

Directory Server now replaces the virtual attribute **nsRole** with an indexed attribute for managed and filtered roles

Previously, LDAP searches that contained the virtual attribute **nsRole** in the filter were time consuming because that attribute cannot be indexed. With this update, when you perform the **ldapsearch** with virtual attribute **nsRole** in the filter, Directory Server replaces the **nsRole** attribute the following way:

- For managed roles, the **nsRole** attribute is replaced with the **nsRoleDN** attribute.
- For filtered roles, the **nsRole** attribute is replaced with the **nsRoleFilter** attribute.

As a result, response time for search with the **nsRole** attribute improves because the search becomes indexed.

Note that this update does not apply to nested roles.

[Bugzilla:2189954](#)

New **nsslapd-numlisteners** configuration option is now available

The **nsslapd-numlisteners** attribute specifies the number of listener threads Directory Server can use to monitor established connections. You can improve the response times when the server experiences a large number of client connection by increasing the attribute value.

[Bugzilla:1975930](#)

IdM supports the option to control the encryption type used to sign the PAC

By default, the Kerberos Key Distribution Center (KDC) generates an AES HMAC-SHA2 signature for the Privilege Attribute Certificate (PAC). However, this encryption type is not supported by Active Directory (AD). As a result, AD cross-realm constrained delegation requests are not processed correctly.

With this enhancement, you can now control the encryption type used to sign the PAC by setting the **pac_privsvr_etype** attribute on the TGS principal, **krbtgt/[realm]@[realm]**, to the required encryption type for the target realm. In IdM, this string attribute is automatically configured when an AD trust exists.

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

For example, for an **MIT** realm and an **AD** realm, to ensure cross-realm ticket-granting tickets (TGT) use AD-compatible encryption types, an administrator must configure the cross-realm TGS principal as shown below on the MIT side. This results in cross-realm TGTs using the AES 256 HMAC-SHA1 encryption type and constrained delegation requests being processed correctly.

```
kadmin.local <<EOF
setstr krbtgt/AD@IPA pac_privsvr_etype aes256-cts-hmac-sha1-96
setstr krbtgt/IPA@AD pac_privsvr_etype aes256-cts-hmac-sha1-96
EOF
```

[Bugzilla:2060421](#)

Identity Management API is now fully supported

The Identity Management (IdM) API was available as a Technology Preview in RHEL 9.2 and as of RHEL 9.3, it is fully supported.

Users can use existing tools and scripts even if the IdM API is enhanced to enable multiple versions of API commands. These enhancements do not change the behavior of a command in an incompatible way. This has the following benefits:

- Administrators can 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.

The communication with the server is possible, regardless if one side uses, for example, a newer version that introduces new options for a feature.

NOTE

While IdM API provides a JSON-RPC interface, this type of access is not supported. Red Hat recommends accessing the API with Python instead. Using Python automates important parts such as the metadata retrieval from the server, which allows listing all available commands.

[Bugzilla:1513934](#)

4.15. GRAPHICS INFRASTRUCTURES

Intel Arc A-Series graphics is now fully supported

The Intel Arc A-Series graphics (Alchemist or DG2) feature, previously available as a Technology Preview, is now fully supported. Intel Arc A-Series graphics is a GPU that enables hardware acceleration, mostly used in PC gaming.

[Bugzilla:2101598^{\[1\]}](#)

4.16. THE WEB CONSOLE

Podman health check action is now available

You can select one of the following Podman health check actions when creating a new container:

- No action (default): Take no action.
- Restart: Restart the container.
- Stop: Stop the container.
- Force stop: Force stops the container, it does not wait for the container to exit.

[Jira:RHELDPCS-16247^{\[1\]}](#)

Stratis is now available in the RHEL web console

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

To learn more about Stratis, see [Setting up Stratis file systems using the web console](#).

[Jira:RHELPLAN-122345^{\[1\]}](#)

4.17. RED HAT ENTERPRISE LINUX SYSTEM ROLES

New RHEL system role for managing `systemd` units

The `rhel-system-role` package now contains the `systemd` RHEL system role. You can use this role to deploy unit files and manage `systemd` units on multiple systems. You can automate `systemd` functionality by providing `systemd` unit files and templates, and by specifying the state of those units, such as started, stopped, masked and other.

[Bugzilla:2224384](#)

New option in the `ssh` role to disable configuration backups

You can now prevent old configuration files from being backed up before they are overwritten by setting the new `ssh_backup` option to `false`. Previously, backup configuration files were created automatically, which might be unnecessary. The default value of the `ssh_backup` option is `true`, which preserves the original behavior.

[Bugzilla:2216753](#)

`keylime_server` RHEL system role

With the new `keylime_server` RHEL system role, you can use Ansible Playbooks to configure the verifier and registrar Keylime components on RHEL 9 systems. Keylime is a remote machine attestation tool that uses the trusted platform module (TPM) technology.

[Bugzilla:2224385](#)

Support for new `ha_cluster` system role features

The `ha_cluster` system role now supports the following features:

- Configuration of resource and resource operation defaults, including multiple sets of defaults with rules.
- Loading and blocking of SBD watchdog kernel modules. This makes installed hardware watchdogs available to the cluster.
- Assignment of distinct passwords to the cluster hosts and the quorum device. This allows you to configure a deployment where the same quorum hosts are joined to multiple, separate clusters, and the passwords of the `hacluster` user on these clusters are different.

For information about the parameters you configure to implement these features, see [Configuring a high-availability cluster by using the `ha_cluster` RHEL system role](#).

[Bugzilla:2185065](#), [Bugzilla:2185067](#), [Bugzilla:2216481](#)

storage system role supports configuring the stripe size for RAID LVM volumes

With this update, you can now specify a custom stripe size when creating RAID LVM devices. For better performance, use the custom stripe size for SAP HANA. The recommended stripe size for RAID LVM volumes is 64 KB.

[Bugzilla:2181656](#)

The `network` RHEL system role supports the `auto-dns` option to control automatic DNS record updates

This enhancement provides support for defined name servers and search domains. You can now use only the name servers and search domains specified in `dns` and `dns_search` properties while disabling automatically configured name servers and search domains such as `dns record` from DHCP. With this enhancement, you can disable automatically auto dns record by changing the `auto-dns` settings.

[Bugzilla:2211194](#)

The `network` RHEL system role supports the `no-aaaa` DNS option

You can now use the `no-aaaa` option to configure DNS settings on managed nodes. Previously, there was no option to suppress AAAA queries generated by the stub resolver, including AAAA lookups triggered by NSS-based interfaces such as `getaddrinfo`; only DNS lookups were affected. With this enhancement, you can now suppress AAAA queries generated by the stub resolver.

[Bugzilla:2218592](#)

The `ad_integration` RHEL system role can now rejoin an AD domain

With this update, you can now use the `ad_integration` RHEL system role to rejoin an Active Directory (AD) domain. To do this, set the `ad_integration_force_rejoin` variable to `true`. If the `realm_list` output shows that host is already in an AD domain, it will leave the existing domain before rejoining it.

[Bugzilla:2211723](#)

The `certificate` RHEL system role now allows changing certificate file mode when using `certmonger`

Previously, certificates created by the **certificate** RHEL system role with the **certmonger** provider used a default file mode. However, in some use-cases you might require a more restrictive mode. With this update, you can now set a different certificate and a key file mode using the **mode** parameter.

[Bugzilla:2218204](#)

The **postgresql** RHEL system role is now available

The new **postgresql** RHEL system role installs, configures, manages, and starts the **PostgreSQL** server. The role also optimizes the database server settings to improve performance.

The role supports the currently released and supported versions of **PostgreSQL** on RHEL 8 and RHEL 9 managed nodes.

For more information, see [Installing and configuring PostgreSQL by using the postgresql RHEL system role](#).

[Bugzilla:2151373](#)

podman RHEL system role now supports Quadlets, health checks, and secrets

Starting with Podman 4.6, you can use the **podman_quadlet_specs** variable in the **podman** RHEL system role. You can define a Quadlet by specifying a unit file, or in the inventory by a name, a type of unit, and a specification. Types of a unit can be the following: **container**, **kube**, **network**, and **volume**. Note that Quadlets work only with root containers on RHEL 8. Quadlets work with rootless containers on RHEL 9.

The health checks are supported only for Quadlet Container types. In the **[Container]** section, specify the **HealthCmd** field to define the health check command and **HealthOnFailure** field to define the action when a container is unhealthy. Possible options are **none**, **kill**, **restart**, and **stop**.

You can use the **podman_secrets** variable to manage secrets. For details, see [upstream documentation](#).

[Jira:RHELPLAN-154441^{\[1\]}](#)

Improved performance of the **selinux** system role with **restorecon -T 0**

The **selinux** system role now uses the **-T 0** option with the **restorecon** command in all applicable cases. This improves the performance of tasks that restore default SELinux security contexts on files.

[Bugzilla:2179460](#)

The **rhc** system role now supports setting a proxy server type

The newly introduced attribute **scheme** under the **rhc_proxy** parameter enables you to configure the proxy server type by using the **rhc** system role. You can set two values: **http**, the default and **https**.

[Bugzilla:2211748](#)

firewall RHEL system role supports variables related to **ipsets**

With this update of the **firewall** RHEL system role, you can define, modify, and delete **ipsets**. Also, you can add and remove those **ipsets** from firewall zones. Alternatively, you can use those **ipsets** when defining firewall rich rules.

You can manage **ipsets** with the **firewall** RHEL system role using the following variables:

- **ipset**
- **ipset_type**
- **ipset_entries**
- **short**
- **description**
- **state: present** or **state: absent**
- **permanent: true**

The following are some notable benefits of this enhancement:

- You can reduce the complexity of the rich rules that define rules for many IP addresses.
- You can add or remove IP addresses from sets as needed without modifying multiple rules.

For more details, see resources in the **/usr/share/doc/rhel-system-roles/firewall/** directory.

[Bugzilla:2229802](#)

RHEL system roles now have new volume options for mount point customization

With this update, you can now specify **mount_user**, **mount_group**, and **mount_permissions** parameters for your mount directory.

[Bugzilla:2181657](#)

The firewall RHEL system role has an option to disable conflicting services, and it no longer fails if firewalld is masked

Previously, the **firewall** system role failed when the **firewalld** service was masked on the role run or in the presence of conflicting services. This update brings two notable enhancements:

The **linux-system-roles.firewall** role always attempts to install, unmask, and enable the **firewalld** service on role run. You can now add a new variable **firewall_disable_conflicting_services** to your playbook to disable known conflicting services, for example, **iptables.service**, **nftables.service**, and **ufw.service**. The **firewall_disable_conflicting_services** variable is set to **false** by default. To disable conflicting services, set the variable to **true**.

[Bugzilla:2222761](#)

Resetting the firewall RHEL system role configuration now requires minimal downtime

Previously, when you reset the **firewall** role configuration by using the **previous: replaced** variable, the **firewalld** service restarted. Restarting adds downtime and prolongs the period of an open connection in which **firewalld** does not block traffic from active connections. With this enhancement, the **firewalld** service completes the configuration reset by reloading instead of restarting. Reloading minimizes the downtime and reduces the opportunity to bypass firewall rules. As a result, using the **previous: replaced** variable to reset the **firewall** role configuration now requires minimal downtime.

[Bugzilla:2223764](#)

4.18. VIRTUALIZATION

sevctl is now fully compatible with AMD EPYC Rome and Milan

With this update, the **sevctl** utility correctly recognizes the latest AMD EPYC cores, including the AMD EPYC Rome and AMD EPYC Milan series. As a result, you can use **sevctl** to configure the features of AMD Secure Encrypted Virtualization (SEV) that are available on these CPUs.

Note, however, that advanced SEV functions, such as SEV-ES and SEV-SNP are only provided as Technology Previews in RHEL 9, and therefore unsupported.

Bugzilla:2104857^[1]

virtio-vga and virtio-gpu devices now support blob resources

It is now possible for **virtio-vga** and **virtio-gpu** devices to use **blob** memory resources, which improves their performance in certain scenarios. To attach a **blob** resource to a **virtio** graphics device, add a **blob="on"** option to the corresponding **<video>** section in the virtual machine's XML configuration. For example:

```
<video>
  <model type="virtio" heads="1" primary="yes" blob="on"/>
  <address type="pci" domain="0x0000" bus="0x00" slot="0x01" function="0x0"/>
</video>
```

Note, however, that this feature currently does not work on IBM Z hosts.

Bugzilla:2032406

Virtualization support for 4th Generation Intel Xeon Scalable processors

With this update, virtualization on RHEL 9 adds support for the 4th Generation Intel Xeon Scalable processors, formerly known as Sapphire Rapids. As a result, virtual machines hosted on RHEL 9 can now use the **SapphireRapids** CPU model and utilise new features that the processors provide.

Bugzilla:1880531^[1]

Improved memory reclaiming for Secure Execution on IBM Z

When using a virtual machine (VM) with IBM Secure Execution on IBM Z, you can now set up enhanced memory reclaiming for the VM. If the VM is using 32 GiB or more RAM, this setting improves the performance of rebooting or stopping the VM.

To set up enhanced memory reclaiming in a VM, add the **<async-teardown enabled='yes'/>** line to the **<features></features>** section in its XML configuration.

Bugzilla:2168499^[1]

New virtualization features in the RHEL web console

With this update, the RHEL web console includes new features in the Virtual Machines page. You can now:

- Select the **Create and edit** button for a virtual machine (VM) based on a cloud image, which allows you to edit all of the VM properties before the VM is installed.
- Create a **raw** storage volume during virtual machine creation.
- Set up a virtual socket (vsock) to enable communication between the host and the VM over a socket.

Note that a virtual socket requires vsock-aware software, such as **socat**, to enable the communication.

Jira:RHELDPCS-16487^[1]

4.19. RHEL IN CLOUD ENVIRONMENTS

cloud-init supports NetworkManager keyfiles

With this update, the **cloud-init** utility can use a NetworkManager (NM) keyfile to configure the network of the created cloud instance.

Note that by default, **cloud-init** still uses the **sysconfig** method for network setup. To configure **cloud-init** to use a NM keyfile instead, edit the `/etc/cloud/cloud.cfg` and set **network-manager** as the primary network renderer:

```
# cat /etc/cloud/cloud.cfg
network:
  renderers: ['network-manager', 'eni', 'netplan', 'sysconfig', 'networkd']
```

Bugzilla:2118235^[1]

cloud-init now uses VMware datasources by default on ESXi

When creating RHEL virtual machines (VMs) on a host that uses the VMware ESXi hypervisor, such as the VMware vSphere cloud platform. This improves the performance and stability of creating an ESXi instance of RHEL by using **cloud-init**. Note, however, that ESXi is still compatible with Open Virtualization Format (OVF) datasources, and you can use an OVF datasource if a VMware one is not available.

Bugzilla:2172341^[1]

4.20. SUPPORTABILITY

sos rebased to version 4.6

The **sos** utility, for collecting configuration, diagnostic, and troubleshooting data, has been rebased to version 4.6. This update provides the following enhancements:

- **sos** reports now include the contents of both `/boot/grub2/custom.cfg` and `/boot/grub2/user.cfg` files that might contain critical information for troubleshooting boot issues. (BZ#2213951)
- The **sos** plugin for OVN-Kubernetes collects additional logs for the interconnect environment. With this update, **sos** also collects logs from the **ovnkube-controller** container when both **ovnkube-node** and **ovnkube-controller** containers are merged into one.

In addition, notable bug fixes include:

- **sos** now correctly gathers **cgroup** data in the OpenShift Container Platform 4 environment (BZ#2186361).
- While collecting **sos** reports with the **sudo** plugin enabled, **sos** now removes the **bindpw** option properly. (BZ#2143272)

- The **subscription_manager** plugin no longer collects proxy usernames and passwords from the `/var/lib/rhsm/` path. (BZ#2177282)
- The **virsh** plugin no longer collects the SPICE remote-display passwords in virt-manager logs, which prevents **sos** from disclosing passwords in its reports. (BZ#2184062)
- **sos** now masks usernames and passwords previously displayed in the `/var/lib/iscsi/nodes/<IQN>/<PortallP>/default` file.



IMPORTANT

The generated archive might contain data considered sensitive. Thus, you should always review the content before passing it to any third party.

(BZ#2187859)

- **sos** completes the tailed log collection even when the size of the log file is exceeded and when a plugin times out. (BZ#2203141)
- When entering the **sos collect** command on a Pacemaker cluster node, **sos** collects an sos report from the same cluster node. (BZ#2186460)
- When collecting data from a host in the OpenShift Container Platform 4 environment, **sos** now uses the **sysroot** path, which ensures that only the correct data are assembled. (BZ#2075720)
- The **sos report --clean** command obfuscates all MAC addresses as intended. (BZ#2207562)
- Disabling the **hpssm** plugin no longer raises exceptions. (BZ#2216608)
- The **sos clean** command follows permissions of sanitized files. (BZ#2218279)

For details on each release of **sos**, see [upstream release notes](#).

Jira:RHELPLAN-156196^[1]

4.21. CONTAINERS

Podman supports pulling and pushing images compressed with zstd

You can pull and push images compressed with the **zstd** format. The zstd compression is more efficient and faster than gzip. It can reduce the amount of network traffic and storage involved in pulling and pushing the image.

Jira:RHELPLAN-154314^[1]

Quadlet in Podman is now available

Beginning with Podman v4.6, you can use Quadlet to automatically generate a **systemd** service file from a container description. The Quadlets might be easier to use than the **podman generate systemd** command because the description focuses on the relevant container details and without the technical complexity of running containers under **systemd**.

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

Jira:RHELPLAN-154432^[1]

The Container Tools packages have been updated

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

Notable changes in Podman v4.6 include:

- The **podman kube play** command now supports the **--configmap=<path>** option to provide Kubernetes YAML file with environment variables used within the containers of the pod.
- The **podman kube play** command now supports multiple Kubernetes YAML files for the **--configmap** option.
- The **podman kube play** command now supports containerPort names and port numbers within liveness probes.
- The **podman kube play** command now adds the ctrName as an alias to the pod network.
- The **podman kube play** and **podman kube generate** commands now support SELinux filetype labels and ulimit annotations.
- A new command, **podman secret exists**, has been added, which verifies if a secret with the given name exists.
- The **podman create**, **podman run**, **podman pod create**, and **podman pod clone** commands now support a new option, **--shm-size-systemd**, which allows limiting tmpfs sizes for systemd-specific mounts.
- The **podman create** and **podman run** commands now support a new option, **--security-opt label=nested**, which allows SELinux labeling within a confined container.
- Podman now supports auto updates for containers running inside a pod.
- Podman can now use an SQLite database as a backend for increased stability. The default remains the BoltDB database. You can select the database by setting the **database_backend** field in the **containers.conf** file.
- Podman now supports Quadlets to automatically generate a **systemd** service file from the container description. The description focuses on the relevant container details and hides the technical complexity of running containers under **systemd**.

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

Jira:RHELPLAN-154438^[1]

Podman now supports a Podmash login shell

Beginning with Podman v4.6, you can use the **Podmash** login shell to manage user access and control. Configure your settings to use the **/usr/bin/podmash** command as a login shell instead of a standard shell command, for example, **/usr/bin/bash**. When a user logs into a system setup, the **podmash** command runs the user's session into a Podman container named **podmash**. Containers into which users log in are defined using the Quadlet files, which are created in the **/etc/containers/systemd/users/** directory. In these files, set the **ContainerName** field in the **[Container]** section to **podmash**. The systemd automatically starts **podmash** when the user session starts and continues running until all user sessions exit.

For more information, see [Podman v4.6.0 Introduces Podmash: A Revolutionary Login Shell](#) .

Jira:RHELPLAN-163003^[1]

Clients for sigstore signatures with Fulcio and Rekor are now available

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, previously available as a Technology Preview, are now fully supported. 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-160660^[1]

The pasta networking mode is now available

Starting with Podman v4.4.1, you can use the **pasta** network mode. It is a high-performance replacement of the default network mode **slirp4netns** and supports IPv6 forwarding. To select the **pasta** network mode, install the **passt** package to use the **podman run** command with the **--network=pasta** option. With Podman v4.6, you can set default rootless network mode in the **/etc/containers/containers.conf** configuration file by using the **default_rootless_network_cmd** field under the **[network]** section.

Jira:RHELDOCS-16240^[1]

UBI 9 Micro Container Image no longer contains **zoneinfo** installed by **tzdata**

With this update, the time zone information provided by the **tzdata** package is no longer included in UBI 9 Micro container images, consequently reducing the image size. The UBI 9 Minimal and UBI 9 Micro containers are UTC-only, and users should reinstall the **tzdata** package to get the full **zoneinfo**, if needed.

[Bugzilla:2223028](#)

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 9.3. These changes could include, for example, added or updated **proc** entries, **sysctl**, and **sysfs** default values, boot parameters, kernel configuration options, or any noticeable behavior changes.

New kernel parameters

amd_pstate=[X86]

With this kernel parameter, you can scale the performance of the AMD CPU. Available values include:

- **disable** - Do not enable **amd_pstate** as the default scaling driver for the supported processors.
- **passive** - Use **amd_pstate** with passive mode as a scaling driver. In this mode autonomous selection is disabled. Driver requests a required performance level and platform tries to match the same performance level if it is satisfied by guaranteed performance level.
- **active** - Use **amd_pstate_epp** driver instance as the scaling driver, driver provides a hint to the hardware if software wants to bias toward performance (0x0) or energy efficiency (0xff) to the CPPC firmware. Then CPPC power algorithm will calculate the runtime workload and adjust the realtime cores frequency.
- **guided** - Activate guided autonomous mode. Driver requests minimum and maximum performance level and the platform autonomously selects a performance level in this range and appropriate to the current workload.

arm64.nosve=[ARM64]

With this kernel parameter, you can unconditionally disable Scalable Vector Extension support.

arm64.nosme=[ARM64]

With this kernel parameter, you can unconditionally disable Scalable Matrix Extension support.

gather_data_sampling=[X86,INTEL]

With this kernel parameter, you can control the Gather Data Sampling (GDS) mitigation. GDS is a hardware vulnerability that allows unprivileged speculative access to data that was previously stored in vector registers.

This issue is mitigated by default in updated microcode. The mitigation might have a performance impact but can be disabled. On systems without the microcode mitigation disabling AVX serves as a mitigation. Available values include:

- **force** - Disable AVX to mitigate systems without microcode mitigation. No effect if the microcode mitigation is present. Known to cause crashes in userspace with buggy AVX enumeration.
- **off** - Disable GDS mitigation.

nospectre_bhb=[ARM64]

With this kernel parameter, you can disable all mitigations for Spectre-BHB (branch history injection) vulnerability. System might allow data leaks with this option.

trace_clock=[FTRACE]

With this kernel parameter, you can set the clock used for tracing events at boot up. Available values include:

- **local** - Use the per CPU timestamp counter.
- **global** - Event timestamps are synchronize across CPUs. Might be slower than the local clock, but better for some race conditions.
- **counter** - Simple counting of events (1, 2, ..) note, some counts might be skipped due to the infrastructure grabbing the clock more than once per event.
- **uptime** - Use jiffies as the timestamp.
- **perf** - Use the same clock that perf uses.
- **mono** - Use the `ktime_get_mono_fast_ns()` function for timestamps.
- **mono_raw** - Use the `ktime_get_raw_fast_ns()` function for timestamps.
- **boot** - Use the `ktime_get_boot_fast_ns()` function for timestamps. Architectures might add more clocks, see [Documentation/trace/ftrace.rst](#) for more details.

Updated kernel parameters

`cgroup.memory=[KNL]`

With this kernel parameter, you can pass options to the **cgroup** memory controller.

- This parameter takes the format of: **<string>**
Available values include:
- **nosocket** - Disable socket memory accounting.
- **nokmem** - Disable kernel memory accounting.
- **[NEW] nobpf** - Disable BPF memory accounting.

`hugetlb_free_vmemmap=[KNL]`


This kernel parameter enables the feature of freeing unused **vmemmap** pages associated with each hugetlb page on boot. For this parameter to work, the

CONFIG_HUGETLB_PAGE_OPTIMIZE_VMEMMAP configuration option must be enabled.

This parameter takes the format of: **{ on | off (default) }**

Available values include:

- **on** - enables this feature
- **off** - disables this feature



NOTE

The **vmemmap** pages might be allocated from the added memory block itself when the **memory_hotplug.memmap_on_memory** module parameter is enabled. Those **vmemmap** pages cannot be optimized even if this feature is enabled. Other **vmemmap** pages not allocated from the added memory block itself are not affected.

intel_pstate=[X86]

You can use this kernel parameter for CPU performance scaling. Available values include:

- **disable** - Do not enable **intel_pstate** as the default scaling driver for the supported processors.
- **[NEW] active** - Use **intel_pstate** driver to bypass the scaling governors layer of **cpufreq** and provides it own algorithms for p-state selection. There are two P-state selection algorithms provided by **intel_pstate** in the active mode: powersave and performance. The way they both operate depends on whether or not the hardware managed P-states (HWP) feature has been enabled in the processor and possibly on the processor model.
- **passive** - Use **intel_pstate** as a scaling driver, but configure it to work with generic **cpufreq** governors (instead of enabling its internal governor). This mode cannot be used along with the hardware-managed P-states (HWP) feature.
- **force** - Enable **intel_pstate** on systems that prohibit it by default in favor of **acpi-cpufreq**. Forcing the **intel_pstate** driver instead of **acpi-cpufreq** might disable platform features, such as thermal controls and power capping, that rely on ACPI P-States information being indicated to OSPM and therefore should be used with caution. This option does not work with processors that are not supported by the **intel_pstate** driver or on platforms that use **pcc-cpufreq** instead of **acpi-cpufreq**.
- **no_hwp** - Do not enable hardware P state control (HWP) if available.
- **hwp_only** - Only load **intel_pstate** on systems that support hardware P state control (HWP) if available.
- **support_acpi_ppc** - Enforce **ACPI_PPC** performance limits. If the Fixed ACPI Description Table specifies preferred power management profile as "Enterprise Server" or "Performance Server", then this feature is turned on by default.
- **per_cpu_perf_limits** - Allow per-logical-CPU P-State performance control limits using the **cpufreq sysfs** interface.

kvm-arm.mode=[KVM,ARM]

With this kernel parameter, you can select one of KVM/arm64's modes of operation. Available values include:

- **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. Setting mode to **protected** disables **kexec** and hibernation for the host.
- **[NEW] nested** - VHE-based mode with support for nested virtualization. Requires at least ARMv8.3 hardware. The **nested** option is experimental and should be used with extreme caution.
Defaults to VHE/nVHE based on hardware support.

libata.force=[LIBATA]

With this kernel parameter, you can force configurations.

The format is a comma-separated list of "[ID:]VAL" where ID is PORT[.DEVICE]. PORT and DEVICE are decimal numbers matching port, link or device. Basically, it matches the ATA ID string printed on console by **libata**.

- If the whole ID part is omitted, the last **PORT** and **DEVICE** values are used.
- If ID has not been specified yet, the configuration applies to all ports, links and devices.
- If only the **DEVICE** value is omitted, the parameter applies to the port and all links and devices behind it. DEVICE number of 0 either selects the first device or the first fan-out link behind PMP device. It does not select the host link. DEVICE number of 15 selects the host link and device attached to it.
- The VAL specifies the configuration to force. As long as there is no ambiguity, shortcut notation is allowed. For example, both 1.5 and 1.5G would work for 1.5Gbps. With the **libata.force=** parameter, you can force the following configurations:
 - Cable type: 40c, 80c, short40c, unk, ign or sata. Any ID with matching PORT is used.
 - SATA link speed limit: 1.5Gbps or 3.0Gbps.
 - Transfer mode: pio[0-7], mwdma[0-4] and udma[0-7]. udma[/][16,25,33,44,66,100,133] notation is also allowed.
 - **nohrst, nosrst, norst**: suppress hard, soft and both resets.
 - **rstone**: only attempt one reset during hot-unplug link recovery.
 - **[NEW] [no]dbdelay**: Enable or disable the extra 200ms delay before debouncing a link PHY and device presence detection.
 - **[no]ncq**: Turn on or off NCQ.
 - **[no]ncqtrim**: Enable or disable queued DSM TRIM.
 - **[NEW] [no]ncqati**: Enable or disable NCQ trim on ATI chipset.
 - **[NEW] [no]trim**: Enable or disable (unqueued) TRIM.
 - **[NEW] trim_zero**: Indicate that TRIM command zeroes data.
 - **[NEW] max_trim_128m**: Set 128M maximum trim size limit.
 - **[NEW] [no]dma**: Turn on or off DMA transfers.
 - **atapi_dmadir**: Enable ATAPI DMADIR bridge support.
 - **atapi_mod16_dma**: Enable the use of ATAPI DMA for commands that are not a multiple of 16 bytes.
 - **[no]dmalog**: Enable or disable the use of the READ LOG DMA EXT command to access logs.
 - **[no]iddevlog**: Enable or disable access to the identify device data log.
 - **[no]logdir**: Enable or disable access to the general purpose log directory.
 - **[NEW] max_sec_128**: Set transfer size limit to 128 sectors.

- **[NEW] max_sec_1024**: Set or clear transfer size limit to 1024 sectors.
- **[NEW] max_sec_lba48**: Set or clear transfer size limit to 65535 sectors.
- **[NEW] [no]lpm**: Enable or disable link power management.
- **[NEW] [no]setxfer**: Indicate if transfer speed mode setting should be skipped.
- **[NEW] [no]fua**: Disable or enable FUA (Force Unit Access) support for devices supporting this feature.
- **dump_id**: Dump IDENTIFY data.
- **disable**: Disable this device.

**NOTE**

If there are multiple matching configurations changing the same attribute, the last one is used.

mitigations=[X86,PPC,S390,ARM64]

With this kernel parameter, you can control optional mitigations for CPU vulnerabilities. This is a set of curated, arch-independent options, each of which is an aggregation of existing arch-specific options. Available values include:

- **off** - disable all optional CPU mitigations. This improves system performance, but it can also expose users to several CPU vulnerabilities. The **off** value is equivalent to:
 - if nokaslr then kpti=0 [ARM64]
 - gather_data_sampling=off [X86]
 - kvm.nx_huge_pages=off [X86]
 - l1tf=off [X86]
 - mds=off [X86]
 - mmio_stale_data=off [X86]
 - no_entry_flush [PPC]
 - no_uaccess_flush [PPC]
 - nobp=0 [S390]
 - nopti [X86,PPC]
 - nospectre_bhb [ARM64]
 - nospectre_v1 [X86,PPC]
 - nospectre_v2 [X86,PPC,S390,ARM64]
 - retbleed=off [X86]

- `spec_store_bypass_disable=off` [X86,PPC]

- `spectre_v2_user=off` [X86]

- `srbds=off` [X86,INTEL]

- `ssbd=force-off` [ARM64]

- `tsx_async_abort=off` [X86]

Exceptions: This does not have any effect on `kvm.nx_huge_pages` when `kvm.nx_huge_pages=force`.

- **auto** (default) - Mitigate all CPU vulnerabilities, but leave SMT enabled, even if it is vulnerable. This is for users who do not want to be surprised by SMT getting disabled across kernel upgrades, or who have other ways of avoiding SMT-based attacks.

- **auto,nosmt** - Mitigate all CPU vulnerabilities, disabling SMT if needed. This is for users who always want to be fully mitigated, even if it means losing SMT. The **auto,nosmt** options are equivalent to:

- `l1tf=flush,nosmt` [X86]

- `mds=full,nosmt` [X86]

- `tsx_async_abort=full,nosmt` [X86]

- `mmio_stale_data=full,nosmt` [X86]

- `retbleed=auto,nosmt` [X86]

nomodeset

With this kernel parameter, you can disable kernel modesetting. Most systems' firmware sets up a display mode and provides framebuffer memory for output. With **nomodeset**, DRM and **fbdev** drivers will not load if they could possibly displace the preinitialized output. Only the system framebuffer will be available for use. The drivers will not perform display-mode changes or accelerated rendering.

This parameter is especially useful as error fallback, or for testing and debugging.

rdt=[HW,X86,RDT]

With this kernel parameter, you can turn on or off individual RDT features. The list includes: **cmt**, **mbmtotal**, **mbmlocal**, **l3cat**, **l3cdp**, **l2cat**, **l2cdp**, **mba**, **smba**, **bmec**.

For example, to turn on **cmt** and turn off **mba** use:

```
rdt=cmt,!mba
```

rodata=[KNL]

With this kernel parameter, you can disable read-only kernel mappings. Available options include:

- **on** - Mark read-only kernel memory as read-only (default).
- **off** - Leave read-only kernel memory writable for debugging.
- **[NEW] full** - Mark read-only kernel memory and aliases as read-only [arm64].

Removed kernel parameters

nobats=[PPC]

With this kernel parameter, you can forbid the use of BATs for mapping kernel lowmem on "Classic" PPC cores.

noltlbs=[PPC]

With this kernel parameter, you can forbid the use of huge page and tlb entries for kernel lowmem mapping on PPC40x and PPC8xx.

swapaccount=[0|1]=[KNL]

With this kernel parameter, you can enable or disable accounting of swap in memory resource controller. For more information, see [Documentation/admin-guide/cgroup-v1/memory.rst](#).

CHAPTER 6. DEVICE DRIVERS

6.1. NEW DRIVERS

Network drivers

- MediaTek MT7601U (USB) support (**mt7601u**), adds support for MT7601U-based wireless USB dongles (only in 64-bit ARM architecture)
- MediaTek MT76x0E (PCIe) support (**mt76x0e**), adds support for MT7610/MT7630-based wireless PCIe devices (only in 64-bit ARM architecture)
- MediaTek MT76x0U (USB) support (**mt76x0u**), adds support for MT7610U-based wireless USB 2.0 dongles (only in 64-bit ARM architecture)
- MediaTek MT76x2E (PCIe) support (**mt76x2e**), adds support for MT7612/MT7602/MT7662-based wireless PCIe devices (only in 64-bit ARM architecture)
- MediaTek MT76x2U (USB) support (**mt76x2u**), adds support for MT7612U-based wireless USB 3.0 dongles (only in 64-bit ARM architecture)
- MediaTek MT7921E (PCIe) support (**mt7921e**), adds support for MT7921E 802.11ax 2x2:2SS wireless devices (only in 64-bit ARM architecture)
- Atheros driver 802.11n HTC based wireless devices (**ath9k_htc**) (only in 64-bit ARM architecture)
- Broadcom 802.11n wireless LAN driver (**brcmsmac**) (only in 64-bit ARM architecture)
- Broadcom 802.11n wireless LAN driver utilities (**brcmutil**) (only in 64-bit ARM architecture)
- Broadcom 802.11 wireless LAN fullmac driver (**brcmfmac**) (only in 64-bit ARM architecture)
- Core module for Qualcomm Atheros 802.11ac wireless LAN cards (**ath10k_core**) (only in 64-bit ARM architecture)
- Core module for Qualcomm Atheros 802.11ax wireless LAN cards (**ath11k**) (only in 64-bit ARM architecture)
- Device simulator for WWAN framework (**wwan_hwsim**)
- Driver support for Qualcomm Atheros 802.11ac WLAN PCIe/AHB devices (**ath10k_pci**) (only in 64-bit ARM architecture)
- Driver support for Qualcomm Technologies 802.11ax WLAN PCIe devices (**ath11k_pci**) (only in 64-bit ARM architecture)
- Intel® Wireless Wi-Fi driver for Linux (**iwlwifi**) (only in 64-bit ARM architecture)
- Intel® Wireless Wi-Fi Link AGN driver for Linux (**iwldvm**)- (only in 64-bit ARM architecture)
- IOSM Driver (**iosm**)
- Marvell WiFi-Ex Driver version 1.0 (**mwifiex**) (only in 64-bit ARM architecture)

- Marvell WiFi-Ex PCI-Express Driver version 1.0 (**mwifiex_pcie**) (only in 64-bit ARM architecture)
- Marvell WiFi-Ex SDIO Driver version 1.0 (**mwifiex_sdio**) (only in 64-bit ARM architecture)
- Marvell WiFi-Ex USB Driver version 1.0 (**mwifiex_usb**) (only in 64-bit ARM architecture)
- MediaTek PCIe 5G WWAN modem T7xx driver (**mtk_t7xx**)
- Network/MBIM over MHI (**mhi_wwan_mbim**) (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures)
- PCI basic driver for rtlwifi (**rtl_pci**) (only in 64-bit ARM architecture)
- Ralink RT2800 library version 2.3.0 (**rt2800lib**) (only in 64-bit ARM architecture)
- Ralink RT2800 PCI & PCMCIA Wireless LAN driver version 2.3.0 (**rt2800pci**) (only in 64-bit ARM architecture)
- Ralink RT2800 USB Wireless LAN driver version 2.3.0 (**rt2800usb**) (only in 64-bit ARM architecture)
- Realtek 802.11ac wireless 8821c driver (**rtw88_8821c**) (only in 64-bit ARM architecture)
- Realtek 802.11ac wireless 8821ce driver (**rtw88_8821ce**) (only in 64-bit ARM architecture)
- Realtek 802.11ac wireless 8822b driver (**rtw88_8822b**) (only in 64-bit ARM architecture)
- Realtek 802.11ac wireless 8822be driver (**rtw88_8822be**) (only in 64-bit ARM architecture)
- Realtek 802.11ac wireless 8822c driver (**rtw88_8822c**) - (only in 64-bit ARM architecture)
- Realtek 802.11ac wireless 8822ce driver (**rtw88_8822ce**) (only in 64-bit ARM architecture)
- Realtek 802.11ac wireless core module (**rtw88_core**) (only in 64-bit ARM architecture)
- Realtek 802.11ac wireless PCI driver (**rtw88_pci**) (only in 64-bit ARM architecture)
- Realtek 802.11ax wireless 8852A driver (**rtw89_8852a**) (only in 64-bit ARM architecture)
- Realtek 802.11ax wireless 8852AE driver (**rtw89_8852ae**) (only in 64-bit ARM architecture)
- Realtek 802.11ax wireless 8852B driver (**rtw89_8852b**) (only in 64-bit ARM architecture and AMD and Intel 64-bit architectures)
- Realtek 802.11ax wireless 8852BE driver (**rtw89_8852be**) (only in 64-bit ARM architecture and AMD and Intel 64-bit architectures)
- Realtek 802.11ax wireless core module (**rtw89_core**) (only in 64-bit ARM architecture)
- Realtek 802.11ax wireless PCI driver (**rtw89_pci**) (only in 64-bit ARM architecture)
- Realtek 802.11n PCI wireless core (**btcoexist**) (only in 64-bit ARM architecture)
- Realtek 802.11n PCI wireless core (**rtlwifi**) (only in 64-bit ARM architecture)
- Realtek 802.11n wireless 8723d driver (**rtw88_8723d**) (only in 64-bit ARM architecture)

- Realtek 802.11n wireless 8723de driver (**rtw88_8723de**) (only in 64-bit ARM architecture)
- Realtek 8188E 802.11n PCI wireless (**rtl8188ee**) (only in 64-bit ARM architecture)
- Realtek 8192C/8188C 802.11n PCI wireless (**rtl8192c-common**) (only in 64-bit ARM architecture)
- Realtek 8192C/8188C 802.11n PCI wireless (**rtl8192ce**) (only in 64-bit ARM architecture)
- Realtek 8192C/8188C 802.11n USB wireless (**rtl8192cu**) (only in 64-bit ARM architecture)
- Realtek 8192DE 802.11n Dual Mac PCI wireless (**rtl8192de**) (only in 64-bit ARM architecture)
- Realtek 8192EE 802.11n PCI wireless (**rtl8192ee**) (only in 64-bit ARM architecture)
- Realtek 8192S/8191S 802.11n PCI wireless (**rtl8192se**) (only in 64-bit ARM architecture)
- Realtek 8723BE 802.11n PCI wireless (**rtl8723be**) (only in 64-bit ARM architecture)
- Realtek 8723E 802.11n PCI wireless (**rtl8723ae**) (only in 64-bit ARM architecture)
- Realtek 8821ae 802.11ac PCI wireless (**rtl8821ae**) (only in 64-bit ARM architecture)
- Realtek RTL8723AE/RTL8723BE 802.11n PCI wireless common routines (**rtl8723-common**) (only in 64-bit ARM architecture)
- rt2800 MMIO library version 2.3.0 (**rt2800mmio**) (only in 64-bit ARM architecture)
- rt2x00 library version 2.3.0 (**rt2x00lib**) (only in 64-bit ARM architecture)
- rt2x00 mmio library version 2.3.0 (**rt2x00mmio**) (only in 64-bit ARM architecture)
- rt2x00 pci library version 2.3.0 (**rt2x00pci**) (only in 64-bit ARM architecture)
- rt2x00 usb library version 2.3.0 (**rt2x00usb**) (only in 64-bit ARM architecture)
- RTL8XXXu USB mac80211 Wireless LAN Driver (**rtl8xxxu**) (only in 64-bit ARM architecture)
- Shared library for Atheros wireless 802.11n LAN cards (**ath9k_common**) (only in 64-bit ARM architecture)
- Shared library for Atheros wireless LAN cards (**ath**) (only in 64-bit ARM architecture)
- Support for Atheros 802.11n wireless LAN cards (**ath9k_hw**) (only in 64-bit ARM architecture)
- Support for Atheros 802.11n wireless LAN cards (**ath9k**) (only in 64-bit ARM architecture)
- The new Intel® wireless AGN driver for Linux (**iwlvmv**) (only in 64-bit ARM architecture)
- Thunderbolt/USB4 network driver (**thunderbolt_net**)
- USB basic driver for rtlwifi (**rtl_usb**) (only in 64-bit ARM architecture)

Graphics drivers and miscellaneous drivers

- Atheros AR30xx firmware driver 1.0 (**ath3k**) (only in 64-bit ARM architecture)
- BlueFRITZ! USB driver version 1.2 (**bfusb**) (only in 64-bit ARM architecture)

- Bluetooth HCI UART driver version 2.3 (**hci_uart**) (only in 64-bit ARM architecture)
- Bluetooth support for Broadcom devices version 0.1 (**btbcm**) (only in 64-bit ARM architecture)
- Bluetooth support for Intel devices version 0.1 (**btintel**) (only in 64-bit ARM architecture)
- Bluetooth support for MediaTek devices version 0.1 (**btmtk**) (only in 64-bit ARM architecture)
- Bluetooth support for Realtek devices version 0.1 (**btrtl**) (only in 64-bit ARM architecture)
- Bluetooth virtual HCI driver version 1.5 (**hci_vhci**) (only in 64-bit ARM architecture)
- Broadcom Blutionium firmware driver version 1.2 (**bcm203x**) (only in 64-bit ARM architecture)
- Digianswer Bluetooth USB driver version 0.11 (**bpa10x**) (only in 64-bit ARM architecture)
- Generic Bluetooth SDIO driver version 0.1 (**btsdio**) (only in 64-bit ARM architecture)
- Generic Bluetooth USB driver version 0.8 (**btusb**) (only in 64-bit ARM architecture)
- Marvell Bluetooth driver version 1.0 (**btmrvl**) (only in 64-bit ARM architecture)
- Marvell BT-over-SDIO driver version 1.0 (**btmrvl_sdio**) (only in 64-bit ARM architecture)
- Linux device driver of the BMC IPMI SSIF interface (**ssif_bmc**) (only in 64-bit ARM architecture)
- vTPM Driver version 0.1 (**tpm_vtpm_proxy**)
- AMD P-state driver Test module (**amd-pstate-ut**) (only in AMD and Intel 64-bit architectures)
- Compute Express Link (CXL) ACPI driver (**cxl_acpi**) (only in 64-bit ARM architecture and AMD and Intel 64-bit architectures)
- Compute Express Link (CXL) core driver (**cxl_core**)
- Compute Express Link (CXL) port driver (**cxl_port**)
- NVIDIA Tegra GPC DMA Controller driver (**tegra186-gpc-dma**) (only in 64-bit ARM architecture)
- DRM Buddy Allocator (**drm_buddy**) (only in 64-bit IBM Z architecture)
- DRM display adapter helper (**drm_display_helper**) (only in 64-bit IBM Z architecture)
- HID driver for EVision devices (**hid-evision**) (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures)
- Texas Instruments INA3221 HWMon Driver (**ina3221**) (only in 64-bit ARM architecture)
- I3C core (**i3c**) (only in 64-bit ARM architecture)
- Silvaco dual-role I3C master driver (**svc-i3c-master**) (only in 64-bit ARM architecture)
- Microsoft Azure Network Adapter IB driver (**mana_ib**) (only in AMD and Intel 64-bit architectures)
- Soft RDMA transport (**rdma_rxe**)

- i.MX8MP interconnect driver – Generic interconnect drivers for i.MX SOCs (**imx8mp-interconnect**) (only in 64-bit ARM architecture)
- Linux USB Video Class (**uvc**) (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures)
- Common memory handling routines for videobuf2 (**videobuf2-memops**) (only in 64-bit ARM architecture)
- Device node registration for cec drivers (**cec**) (only in 64-bit IBM Z architecture)
- Device node registration for media drivers (**mc**) (only in 64-bit ARM architecture)
- Driver helper framework for Video for Linux 2 (**videobuf2-v4l2**) (only in 64-bit ARM architecture)
- Media buffer core framework (**videobuf2-common**) (only in 64-bit ARM architecture)
- USB Video Class driver version 1.1.1 (**uvcvideo**) (only in 64-bit ARM architecture)
- V4L2 DV Timings Helper Functions (**v4l2-dv-timings**) (only in 64-bit ARM architecture)
- Video4Linux2 core driver (**videodev**) (only in 64-bit ARM architecture)
- vmalloc memory handling routines for videobuf2 (**videobuf2-vmalloc**) (only in 64-bit ARM architecture)
- Framework for SPI NOR (**spi-nor**) (only in 64-bit ARM architecture)
- Marvell CN10K DRAM Subsystem(DSS) PMU (**marvell_cn10k_ddr_pmu**) (only in 64-bit ARM architecture)
- Marvell CN10K LLC-TAD Perf driver (**marvell_cn10k_tad_pmu**) (only in 64-bit ARM architecture)
- Intel Meteor Lake PCH pinctrl/GPIO driver (**pinctrl-meteorlake**) (only in AMD and Intel 64-bit architectures)
- Intel In Field Scan (IFS) device (**intel_ifs**) (only in AMD and Intel 64-bit architectures)
- NVIDIA WMI EC Backlight driver (**nvidia-wmi-ec-backlight**) (only in AMD and Intel 64-bit architectures)
- QMI encoder/decoder helper (**qmi_helpers**) (only in 64-bit ARM architecture)
- AMD SoundWire driver (**soundwire-amd**) (only in AMD and Intel 64-bit architectures)
- NVIDIA Tegra114 SPI Controller Driver (**spi-tegra114**) (only in 64-bit ARM architecture)
- STMicroelectronics STUSB160x Type-C controller driver (**stusb160x**) (only in 64-bit ARM architecture)
- MLX5 VFIO PCI - User Level meta-driver for MLX5 device family (**mlx5-vfio-pci**)

6.2. UPDATED DRIVERS

Network driver updates

- Realtek RTL8152/RTL8153 Based USB Ethernet Adapters (**r8152**) has been updated to version v1.12.13 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).

Storage driver updates

- Broadcom MegaRAID SAS Driver (**megaraid_sas**) has been updated to version 07.725.01.00-rc1 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures).
- Driver for Microchip Smart Family Controller (**smartpqi**) has been updated to version 2.1.22-040 (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 0:14.2.0.12 (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 8.4.1.0.0.

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	528482304
CONFIG_BPF	y
CONFIG_BPF_SYSCALL	y
CONFIG_HAVE_EBPF_JIT	y
CONFIG_BPF_JIT	y
CONFIG_BPF_JIT_ALWAYS_ON	y
CONFIG_DEBUG_INFO_BTF	y
CONFIG_DEBUG_INFO_BTF_MODULES	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	y
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_strotol, 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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_strotol, bpf_strtoul, 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_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_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_copy_from_user_task, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

Program type	Available helpers
sched_cls	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_l3_csum_replace, bpf_l4_csum_replace, bpf_tail_call, bpf_clone_redirect, bpf_get_cgroup_classid, bpf_skb_vlan_push, bpf_skb_vlan_pop, bpf_skb_get_tunnel_key, bpf_skb_set_tunnel_key, bpf_redirect, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_get_tunnel_opt, bpf_skb_set_tunnel_opt, bpf_skb_change_proto, bpf_skb_change_type, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_change_tail, bpf_skb_pull_data, bpf_csum_update, bpf_set_hash_invalid, bpf_get_numa_node_id, bpf_skb_change_head, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_set_hash, bpf_skb_adjust_room, bpf_skb_get_xfrm_state, bpf_skb_load_bytes_relative, bpf_fib_lookup, bpf_skb_cgroup_id, bpf_skb_ancestor_cgroup_id, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_fullsock, bpf_tcp_sock, bpf_skb_ecn_set_ce, bpf_get_listener_sock, bpf_skc_lookup_tcp, bpf_tcp_check_syncookie, bpf_strtol, bpf_strtoul, 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_tstamp, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_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_strtol, bpf_strtoul, 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_strcmp, bpf_skb_set_tstamp, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

Program type	Available helpers
tracepoint	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_perf_event_read_value, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_strotol, bpf_strtoull, 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_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_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_copy_from_user_task, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_strotol, bpf_strtoull, bpf_tcp_gen_syncookie, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_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_xdp_get_buff_len, bpf_xdp_load_bytes, bpf_xdp_store_bytes, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_strotol, bpf_strtoul, 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_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_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_copy_from_user_task, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_strotol, bpf_strtoul, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_sk_cgroup_id, bpf_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_strtol, bpf_strtoul, 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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_strtol, bpf_strtoul, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_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_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_strtol, bpf_strtoul, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_strtol, bpf_strtoul, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_strotol, bpf_strtoul, 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_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_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_copy_from_user_task, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_strotol, bpf_strtoul, 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_get_retval, bpf_set_retval, 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_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_lwt_seg6_store_bytes, bpf_lwt_seg6_adjust_srh, bpf_lwt_seg6_action, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, 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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_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_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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

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_strotol, bpf_strtoul, 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_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_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_copy_from_user_task, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_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_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_strotol, bpf_strtoul, 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_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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_strtol, bpf_strtoul, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
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_strtol, bpf_strtoul, 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, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

Table 7.3. Available map types

Map type	Available
hash	yes
array	yes
prog_array	yes
perf_event_array	yes
percpu_hash	yes
percpu_array	yes
stack_trace	yes
cgroup_array	yes
lru_hash	yes
lru_percpu_hash	yes
lpm_trie	yes
array_of_maps	yes
hash_of_maps	yes
devmap	yes
sockmap	yes
cpumap	yes
xskmap	yes
sockhash	yes
cgroup_storage	yes
reuseport_sockarray	yes
percpu_cgroup_storage	yes
queue	yes
stack	yes

Map type	Available
sk_storage	yes
devmap_hash	yes
struct_ops	yes
ringbuf	yes
inode_storage	yes
task_storage	yes
bloom_filter	yes
user_ringbuf	yes
cgrp_storage	yes

CHAPTER 8. BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 9.3 that have a significant impact on users.

8.1. INSTALLER AND IMAGE CREATION

The installation program now correctly processes the `--proxy` option of the `url` Kickstart command

Previously, the installation program did not correctly process the `--proxy` option of the `url` Kickstart command. As a consequence, you could not use the specified proxy to fetch the installation image. With this update, the issue is fixed and the `--proxy` option now works as expected.

[Bugzilla:2177219](#)

The `--noverifyssl` option for `liveimg` no longer checks the server's certificate for images downloaded using HTTPS

Previously, the installation program ignored the `--noverifyssl` option from the `liveimg` Kickstart command. Consequently, if the server's certificate could not be validated for images downloaded using the HTTPS protocol, the installation process failed. With this update, this issue has been fixed, and the `--noverifyssl` option of the `liveimg` Kickstart command works correctly.

[Bugzilla:2157921](#)

Anaconda now validates LUKS passphrases for the FIPS requirements

Previously, Anaconda did not check whether the length of LUKS passphrases satisfied the FIPS requirements, even though the underlying tools performed this check. As a consequence, installing in FIPS mode with a passphrase shorter than 8 characters caused the installer to stop prematurely.

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

[Bugzilla:2163497](#)

The new version of `xfstools` no longer shrinks the size of `/boot`

Previously, the `xfstools` package with the 5.19 version in the RHEL 9.3 caused the size of `/boot` to shrink. As a consequence, it caused a difference in the available space on the `/boot` partition, if compared to the RHEL 9.2 version. This fix increases the `/boot` partition to 600 MiB for all images, instead of 500 MiB, and the `/boot` partition is no longer affected by space issues.

[Jira:RHEL-7999](#)

8.2. SECURITY

OpenSSL commands `cms` and `smime` can encrypt files in FIPS mode

Previously, the default configuration of the `cms` and `smime` OpenSSL commands used legacy encryption algorithms, such as 3DES or PKCS #1 v1.5. These algorithms are disabled in FIPS mode. As a result, encrypting files by using the `smime` command with the default settings did not work on systems in FIPS mode. This update introduces the following changes:

- In FIPS mode, OpenSSL APIs create CMS data by using OAEP with RSA keys by default.

- In FIPS mode, the **cms** OpenSSL command creates CMS files encrypted with **aes-128-cbc** and OAEP when provided RSA keys.

The use of ECDSA keys is unaffected. In non-FIPS mode, OpenSSL APIs and the **cms** command continue to use PKCS#1 v1.5 padding and 3DES encryption by default.

As a consequence, you can use the **cms** and **smime** OpenSSL commands in FIPS mode to encrypt files.

[Bugzilla:2160797](#)

SELinux allows mail replication in Dovecot

You can configure the Dovecot high-performance mail delivery agent for high availability with two-way replication set, but the SELinux policy previously did not contain rules for the **dovecot-deliver** utility to communicate over a pipe in the runtime filesystem. As a consequence, mail replication in Dovecot did not work. With this update, permissions have been added to the SELinux policy, and as a result, mail replication in Dovecot works.

[Bugzilla:2170495](#)^[1]

Booting from an NFS filesystem now works with SELinux set to enforcing mode

Previously, when using NFS as the root filesystem, SELinux labels were not forwarded from the server, causing boot failures when SELinux was set to enforcing mode.

With this fix, SELinux has been fixed to correctly flag NFS mounts created before the initial SELinux policy load as supporting security labels. As a result, the NFS mount now forwards SELinux labels between the server and the client and the boot can succeed with SELinux set to enforcing mode.

[Bugzilla:2218207](#)^[1]

rabbitmq no longer fails with IPv6

Previously, when you deployed **rabbitmq** server with IPv6 enabled, the **inet_gethost** command tried to access the `/proc/sys/net/ipv6/conf/all/disable_ipv6` file. Consequently, the system denied access to `/proc/sys/net/ipv6/conf/all/disable_ipv6`. With this update, system can now read `/proc/sys/net/ipv6/conf/all/disable_ipv6`, and **rabbitmq** now works with IPv6.

[Bugzilla:2184999](#)

Registration to Insights through cloud-init is no longer blocked by SELinux

Previously, the SELinux policy did not contain a rule that allows the **cloud-init** script to run the **insights-client** service. Consequently, an attempt to run the **insights-client --register** command by the **cloud-init** script failed. With this update, the missing rule has been added to the policy, and you can register to Insights through **cloud-init** with SELinux in enforcing mode.

[Bugzilla:2162663](#)

Users in the staff_r SELinux role can now run scap_workbench probes

Previously, the **selinux-policy** packages did not contain rules for users in the **staff_r** SELinux role required to run the **scap-workbench** utility. Consequently, **scap-workbench** probes failed when run by user in the **staff_r** SELinux role. With this update, the missing rules have been added to **selinux-policy**, and SELinux users can now run **scap_workbench** probes.

[Bugzilla:2112729](#)

Permissions for `insights-client` added to the SELinux policy

The `insights-client` service requires permissions that were not in the previous versions of the `selinux-policy`. As a consequence, some components of `insights-client` did not work correctly and reported access vector cache (AVC) error messages. This update adds new permissions to the SELinux policy. As a result, `insights-client` runs correctly without reporting AVC errors.

Jira:RHELPLAN-163014^[1], Bugzilla:2190178, Bugzilla:2224737, Bugzilla:2207894, Bugzilla:2214581

Keylime allowlist generation script updated

The Keylime script `create_allowlist.sh` generates an allowlist for the Keylime policy. In RHEL 9.3, it was replaced with the `create_runtime_policy.sh` script, which failed when trying to convert the allowlist to the JSON runtime policy.

With this update, the script was reverted to `create_allowlist.sh`. Now, you can combine the allowlist and excludelist into the JSON runtime policy by using the `keylime_create_policy` script.

Jira:RHEL-11866^[1]

Keylime no longer requires a specific file for `tls_dir = default`

Previously, when the `tls_dir` variable was set to `default` in Keylime verifier or registrar configuration, Keylime rejected custom certificate authority (CA) certificates that had a different file name than `cacert.crt`. With this update, the problem no longer occurs, and you can use custom CA certificate files even with the `tls_dir = default` setting.

Jira:RHELPLAN-157337^[1]

Environment variables can override Keylime agent options with underscores

Previously, when a Keylime agent configuration option name contained an underscore (`_`), overriding this option through environment variables did not work. With this update, the override through environment variables works correctly even when an option name contains an underscore.

Jira:RHEL-395^[1]

Keylime registrar correctly identifies IPv6 addresses

Previously, the Keylime registrar did not correctly recognize IPv6 addresses, and therefore failed to bind its listening port. With this update, the registrar properly identifies IPv6 addresses and, consequently, binds to its port correctly.

Jira:RHEL-392^[1]

Keylime agent correctly handles IPv6 addresses

Previously, when registering a Keylime agent by using an IPv6 address not enclosed in brackets, `[]`, the `keylime_tenant` utility failed with an error. With this update, `keylime_tenant` handles IPv6 addresses correctly even when they are not enclosed in brackets.

Jira:RHEL-393^[1]

Keylime no longer fails measured boot attestation due to new events in QEMU VMs

An update of the `edk2-ovmf` package introduced a new type of events in the measured boot log for virtual systems operated by QEMU. These events caused failures in Keylime measured boot attestation. With this update, Keylime handles these events correctly.

[Jira:RHEL-947^{\[1\]}](#)

Keylime webhook notifier correctly closes TLS sessions

Previously, the keylime webhook notifier did not correctly close TLS sessions. This caused warnings being reported on the listener side. This update fixed this issue, and the webhook notifier now correctly closes TLS sessions.

[Jira:RHEL-1252^{\[1\]}](#)

gpg-agent now works as an SSH agent in FIPS mode

Previously, the **gpg-agent** tool created MD5 fingerprints when adding keys to the **ssh-agent** program even though FIPS mode disabled the MD5 digest. As a consequence, the **ssh-add** utility failed to add the keys to the authentication agent.

With this release, **gpg-agent** no longer use MD5 checksums. As a result, **gpg-agent** now works as an SSH authentication agent also on systems running in FIPS mode.

[Bugzilla:2073567](#)

tangd-keygen now handles non-default umask correctly

Previously, the **tangd-keygen** script did 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 returned the error message **Internal Error 500** instead of displaying the keys. With this update, **tangd-keygen** sets file permissions for generated key files, and therefore the script now works correctly on systems with non-default **umask**.

[Bugzilla:2188743](#)

fapolicyd service no longer runs programs that are removed from the trusted database

Previously, the **fapolicyd** service incorrectly handled a program as trusted even after it was removed from the trusted database. As a result, entering the **fapolicyd-cli --update** command had no effect, and the program could be executed even after being removed. With this update, the **fapolicyd-cli --update** command correctly updates the trusted programs database, and removed programs can no longer be executed.

[Jira:RHEL-622](#)

fapolicyd no longer causes the system to hang after mount and umount

Previously, when the **mount** or **umount** actions were run twice followed by the **fapolicyd-cli --update** command, the **fapolicyd** service might enter an endless loop. As a result, the system stopped responding. With this update, the service runs the **fapolicyd-cli --update** command correctly, and the service handles any number of **mount** or **umount** actions.

[Jira:RHEL-817](#)

Keylime now accepts concatenated PEM certificates

Previously, when Keylime received a certificate chain as multiple certificates in the PEM format concatenated in a single file, the **keylime-agent-rust** Keylime component produced a TLS handshake failure. As a consequence, the client components (**keylime_verifier** and **keylime_tenant**) could not connect to the Keylime agent. With this update, **keylime-agent-rust** correctly handles multiple certificates including intermediary CA certificates. As a result, you can now use concatenated PEM certificates with Keylime.

Jira:RHEL-396^[1]

Rsyslog can start even without capabilities

When Rsyslog is executed as a normal user or in a containerized environment, the **rsyslog** process has no capabilities. Consequently, Rsyslog in this scenario could not drop capabilities and exited at startup. With this update, the process no longer attempts to drop capabilities if it has no capabilities. As a result, Rsyslog can start even when it has no capabilities.

Jira:RHELPLAN-160541^[1]

io_uring now works without SELinux denials

Previously, the **io_uring** kernel interface missed the **map** permission in the SELinux policy. Consequently, the **mmap** system call failed and the **io_uring** interface did not work properly. With this update, the **map** permissions have been allowed in SELinux policy and the interface now works without SELinux denials.

[Bugzilla:2187745](#)

oscap-anaconda-addon can now harden Network Servers for CIS

Previously, installing RHEL Network Servers with a CIS security profile (**cis**, **cis_server_I1**, **cis_workstation_I1**, or **cis_workstation_I2**) was not possible with the Network Servers package group selected. This problem is fixed by excluding the **tftp** package in **oscap-anaconda-addon-2.0.0-17.el9** provided with RHEL 9.3. As a consequence, you can install CIS-hardened RHEL Network Servers with the Network Servers package group.

[Bugzilla:2172264](#)

Rules checking home directories apply only to local users

Multiple compliance profiles provided by the **scap-security-guide** package contain the following rules that check the correct configuration of user home directories:

- **accounts_umask_interactive_users**
- **accounts_user_dot_group_ownership**
- **accounts_user_dot_user_ownership**
- **accounts_user_interactive_home_directory_exists**
- **accounts_users_home_files_groupownership**
- **accounts_users_home_files_ownership**
- **accounts_users_home_files_permissions**
- **file_groupownership_home_directories**
- **file_ownership_home_directories**
- **file_permissions_home_directories**

These rules correctly check the configuration of local users. Previously, the scanner also incorrectly checked the configuration of remote users provided by network sources such as NSS even though the remediation scripts could not change remote users' configuration. This was because the OpenSCAP

scanner previously used the **getpwent()** system call. This update changes the internal implementation of these rules to depend only on the data from the **/etc/passwd** file. As a result, the rules now apply only to the local users' configuration.

[Bugzilla:2203791](#)

Password age rules apply only to local users

Some compliance profiles, for example CIS and DISA STIG, contain the following rules checking password age and password expiration of user account passwords:

- **accounts_password_set_max_life_existing**
- **accounts_password_set_min_life_existing**
- **accounts_password_set_warn_age_existing**
- **accounts_set_post_pw_existing**

These rules correctly check the configuration of local users. Previously, the scanner also incorrectly checked the configuration of remote users provided by network sources such as NSS even though the remediation scripts could not change remote users' configuration. This was because the OpenSCAP scanner previously used the **getpwent()** system call.

This update changes the internal implementation of these rules to depend only on the data from the **/etc/shadow** file. As a result, the rules now apply only to the local users' configuration.

[Bugzilla:2213958](#)

Red Hat CVE feeds have been updated

The version 1 of Red Hat Common Vulnerabilities and Exposures (CVE) feeds at <https://access.redhat.com/security/data/oval/> has been discontinued and replaced by the version 2 of the CVE feeds located at <https://access.redhat.com/security/data/oval/v2/>.

Consequently, the links in SCAP source data streams provided by the **scap-security-guide** package have been updated to link to the new version of the Red Hat CVE feeds.

[Bugzilla:2223178](#)

Rules related to **journald** configuration no longer add extra quotes

Previously, the SCAP Security Guide rules **journald_compress**, **journald_forward_to_syslog**, and **journald_storage** previously contained a bug in the remediation script which caused adding extra quotes to the configuration options in the **/etc/systemd/journald.conf** configuration file. Consequently, the **journald** system service failed to parse the configuration options and ignored them. Therefore, the configuration options were not effective. This caused false **pass** results in OpenSCAP scans. With this update, the rules and remediations scripts no longer add the extra quotes. As a result, these rules now produce a valid configuration for **journald**.

[Bugzilla:2193169](#)

Files under **/var/lib/fdo** now get the correct SELinux label

Previously, there was a security issue that allowed the FDO process to access the entire host. With this update, by using the **service-info-api** server with SELinux, you can add any file to send to the device under the **/var/lib/fdo** directory, and, as a consequence, the files under **/var/lib/fdo** will now get the correct SELinux label.

[Bugzilla:2229722](#)

8.3. SUBSCRIPTION MANAGEMENT

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^{\[1\]}](#)

8.4. SOFTWARE MANAGEMENT

The dnf needs-restarting -s command now correctly displays the list of systemd services

Previously, when you used the **needs-restarting** command with the **-s** or **--services** option, an error occurred when a non-systemd or malfunctioning process was detected. With this update, the **dnf needs-restarting -s** command ignores such processes and displays a warning instead with the list of affected systemd services.

[Bugzilla:2203100](#)

The dnf-automatic command now correctly reports the exit status of transactions

Previously, the **dnf-automatic** command returned a successful exit code of a transaction even if some actions during this transaction were not successfully completed. This could cause a security risk on machines that use **dnf-automatic** for automatic deployment of errata. With this update, the issue has been fixed and **dnf-automatic** now reports every problem with packages during the transaction.

[Bugzilla:2212262](#)

Installing packages with IMA signatures on file systems without extended file attributes no longer fails

Previously, RPM tried to apply IMA signatures to files even if they did not support these signatures. As a consequence, package installation failed. With this update, RPM skips applying IMA signatures. As a result, package installation no longer fails.

[Bugzilla:2157836](#)

8.5. SHELLS AND COMMAND-LINE TOOLS

The rsyslog logging service now starts at boot of the rescue system

Previously, the **rsyslog** service for message logging did not automatically start in the rescue system. The **/dev/log** socket kept receiving messages during the recovery process with no service listening at this socket. Consequently, the **/dev/log** socket was filled with messages and caused the recovery process to be stuck. For example, the **grub2-mkconfig** command to regenerate the GRUB

configuration produces a high amount of log messages depending on the number of mounted file systems. If you used ReaR to recover systems with many mounted file systems, numerous log messages would fill the `/dev/log` socket, and the recovery process froze.

With this fix, the **systemd** units in the rescue system now include the sockets target in the boot procedure to start the logging socket at boot. As a result, the **rsyslog** service starts in the rescue environment when required, and the processes that need to log messages during recovery are no longer stuck. The recovery process completes successfully and you can find the log messages in the `/var/log/messages` file in the rescue RAM disk.

[Bugzilla:2172912](#)

The **which** command no longer fails for a long path

Previously, when you executed the **which** command in a directory with a path longer than 256 characters, the command failed with the **Can't get current working directory** error message. With this fix, the **which** command now uses the **PATH_MAX** value for the path length limit. As a result, the command no longer fails.

[Bugzilla:2181974](#)

ReaR now supports UEFI Secure Boot with **OUTPUT=USB**

Previously, the **OUTPUT=USB** ReaR output method, which stores the rescue image on a bootable disk drive, did not respect the **SECURE_BOOT_BOOTLOADER** setting. Consequently, on systems with UEFI Secure Boot enabled, the disk with the rescue image would not boot because the boot loader was not signed.

With this fix, the **OUTPUT=USB** ReaR output method now uses the boot loader that you specify in the **SECURE_BOOT_BOOTLOADER** setting when creating the rescue disk. To use the signed UEFI shim boot loader, change the following setting in the `/etc/rear/local.conf` file:

```
SECURE_BOOT_BOOTLOADER=/boot/efi/EFI/redhat/shimx64.efi
```

As a result, the rescue disk is bootable when UEFI Secure Boot is enabled. It is safe to set the variable to this value on all systems with UEFI, even when Secure Boot is not enabled. It is even recommended for consistency. For details about the UEFI boot procedure and the shim boot loader, see [UEFI: what happens when booting the system](#).

[Bugzilla:2196445](#)

System recovered by ReaR no longer fails to mount all VG logical volumes

The `/etc/lvm/devices/system.devices` file represents the Logical Volume Manager (LVM) system devices and controls device visibility and usability to LVM. By default, the **system.devices** feature is enabled in RHEL 9 and when active, it replaces the LVM device filter.

Previously, when you used ReaR to recover the systems to disks with hardware IDs different from those the original system used, the recovered system did not find all LVM volumes and failed to boot. With this fix, if ReaR finds the **system.devices** file, ReaR moves this file to `/etc/lvm/devices/system.devices.rearbak` at the end of recovery. As a result, the recovered system does not use the LVM devices file to restrict device visibility and the system finds the restored volumes at boot.

Optional: If you want to restore the default behavior and regenerate the LVM devices file, use the **vgimportdevices -a** command after booting the recovered system and connecting all disk devices needed for a normal operation, in case you disconnected any disks before the recovery process.

[Bugzilla:2145014](#)

8.6. NETWORKING

Intel Corporation I350 Gigabit Fiber Network Connection now provides a link after kernel update

Previously, hardware configurations with Small Formfactor Pluggable (SFP) transceiver modules without External Thermal Sensor (ETS) caused the **igb** driver to erroneously initialize the Inter-Integrated Circuit (I2C) to read ETS. As a consequence, connections did not obtain links. With this bug fix, the **igb** driver only initializes I2C when SFP with ETS is available. As a result, connections obtain links.

[Bugzilla:2173594^{\[1\]}](#)

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

Based on the information received from the cloud environment, the **nm-cloud-setup** service configured network interfaces. While you had the option to disable **nm-cloud-setup** for manual interface configuration, certain scenarios led to conflicts. In some cases, other services on the host would independently configure interfaces, including the addition of secondary IP addresses. **nm-cloud-setup** incorrectly removed these secondary IP addresses when triggered again by the **systemd** timer unit. This update for the **NetworkManager** package fixes the problem. You only need to wait for the **systemd** timer unit to trigger **nm-cloud-setup**. If you do not want to wait for the timer, you can enable **nm-cloud-setup** manually with the following command:

```
# systemctl enable nm-cloud-setup.service
```

As a result, **nm-cloud-setup** no longer removes manually-configured secondary IP addresses from interfaces.

[Bugzilla:2151040](#)

8.7. KERNEL

RHEL previously failed to recognize NVMe disks when VMD was enabled

When you reset or reattached a driver, the Volume Management Device (VMD) domain previously did not soft-reset. Consequently, the hardware could not properly detect and enumerate its devices. With this update, the operating system with VMD enabled now correctly recognizes NVMe disks, especially when resetting a server or working with a VM machine.

[Bugzilla:2128610^{\[1\]}](#)

8.8. BOOT LOADER

GRUB now correctly handles non-debug kernel variants

Previously, in systems with multiple kernel RPMs installed, entering the **dnf install kernel-\$VERSION** or **dnf update** commands set the last-installed kernel as the default kernel. This occurred, for example, in systems with the standard kernel and real-time kernel on AMD and Intel 64-bit architectures, or kernel (4k) and **kernel-64k** on 64-bit ARM architecture. As a consequence, the system could boot into the unneeded kernel on future reboots. With this update, GRUB uses the **DEFAULTKERNEL** variable in the **/etc/sysconfig/kernel** configuration file, and the default kernel remains the proper variant and latest version.

For more information, see the [Changing the default kernel in Red Hat Enterprise Linux 8 & 9](#) solution.

Bugzilla:2184069^[1]

8.9. FILE SYSTEMS AND STORAGE

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^[1]

`multipathd` adds the persistent reservation registration key to all paths

Previously, when the `multipathd` daemon started and it recognized a registration key for the persistent reservations on one path of an existing multipath device, not all paths of that device had the registration key. As a consequence, if new paths appeared to a multipath device with persistent reservations while `multipathd` was stopped, persistent reservations were not set up on those. This allowed IO processing on the paths, even if they were supposed to be forbidden by the reservation key.

With this fix, if `multipathd` finds a persistent reservation registration key on any device path, it adds the key to all active paths. As a result, multipath devices now have persistent reservations set up correctly on all the paths, even if path devices first appear while `multipathd` is not running.

Bugzilla:2164869

LUNs are now visible during the operating system installation

Previously, the system was not using the authentication information from firmware sources, specifically in cases involving iSCSI hardware offload with CHAP (Challenge-Handshake Authentication Protocol) authentication stored in the iSCSI iBFT (Boot Firmware Table). As a consequence, the iSCSI login failed during installation.

With the fix in the `udisks2-2.9.4-9.el9` firmware authentication, this issue is now resolved and LUNs are visible during the installation and initial boot.

Bugzilla:2213769^[1]

System boots correctly when adding a NVMe-FC device as a mount point in `/etc/fstab`

Previously, due to a known issue in the `nvme-cli nvmf-autoconnect systemd` services, systems failed to boot while adding the Non-volatile Memory Express over Fibre Channel (NVMe-FC) devices as a mount point in the `/etc/fstab` file. Consequently, the system entered into an emergency mode. With this update, a system boots without any issue when mounting an NVMe-FC device.

Jira:RHEL-8171^[1]

8.10. HIGH AVAILABILITY AND CLUSTERS

The `pcs config checkpoint diff` command now works correctly for all configuration sections

As of the RHEL 9.0 release, the `pcs config checkpoint diff` command had stopped showing the differences for the following configuration sections: Fencing Levels, Ordering Constraints, Colocation Constraints, Ticket Constraints, Resources Defaults, and Operations Defaults. As of the RHEL 9.1 release, the `pcs config checkpoint diff` command had stopped showing the differences for the Resources and Stonith devices configuration sections. This is because as the code responsible for displaying each of the different configuration sections switched to a new mechanism for loading CIB files, the loaded content was cached. The second file used for the difference comparison was not loaded and the cached content of the first file was used instead. As a result, the `diff` command yielded no output. With this fix, the CIB file content is no longer cached and the `pcs config checkpoint diff` command shows differences for all configuration sections.

[Bugzilla:2175881](#)

`pcsd` Web UI now displays cluster status when fence levels are configured

Previously, the `pcsd` Web UI did not display cluster status when fence levels were configured. With this fix, you can now view the cluster status and change the cluster settings with the Web UI when fence levels are configured.

[Bugzilla:2182810](#)

A fence watchdog configured as a second fencing device now fences a node when the first device times out

Previously, when a watchdog fencing device was configured as the second device in a fencing topology, the watchdog timeout would not be considered when calculating the timeout for the fencing operation. As a result, if the first device timed out the fencing operation would time out even though the watchdog would fence the node. With this fix, the watchdog timeout is included in the fencing operation timeout and the fencing operation succeeds if the first device times out.

[Bugzilla:2182482](#)

Location constraints with rules no longer displayed when listing is grouped by nodes

Location constraints with rules cannot have a node assigned. Previously, when you grouped the listing by nodes, location constraints with rules were displayed under an empty node. With this fix, the location constraints with rules are no longer displayed and a warning is given indicating that constraints with rules are not displayed.

[Bugzilla:1423473](#)

`pcs` command to update multipath SCSI devices now works correctly

Due to changes in the Pacemaker CIB file, the `pcs stonith update-scsi-devices` command stopped working as designed, causing an unwanted restart of some cluster resources. With this fix, this command works correctly and updates SCSI devices without requiring a restart of other cluster resources running on the same node.

[Bugzilla:2177996](#)

Memory footprint of `pcsd-ruby` daemon now reduced when `pcsd` Web UI is open

Previously, when the `pcsd` Web UI was open, memory usage of the `pcsd-ruby` daemon increased steadily over the course of several hours. With this fix, the web server that runs in the `pcsd-ruby` daemon now periodically performs a graceful restart. This frees the allocated memory and reduces the memory footprint.

[Bugzilla:1860626^{\[1\]}](#)

The **azure-events-az** resource agent no longer produces an error with Pacemaker 2.1 and later

The **azure-events-az** resource agent executes the **crm_simulate -Ls** command and parses the output. With Pacemaker 2.1 and later, the output of the **crm_simulate** command no longer contains the text **Transition Summary:**, which resulted in an error. With this fix, the agent no longer yields an error when this text is missing.

[Bugzilla:2182415](#)

The **mysql** resource agent now works correctly with promotable clone resources

Previously, the **mysql** resource agent moved cloned resources that were operating in a Promoted role between nodes, due to promotion scores changing between promoted and non-promoted values. With this fix, a node in a Promoted role remains in a Promoted role.

[Bugzilla:2179003^{\[1\]}](#)

The **fence_scsi** agent is now able to auto-detect shared **lvmlckd** devices

Previously, the **fence_scsi** agent did not auto-detect shared **lvmlckd** devices. With this update, **fence_scsi** is able to auto-detect **lvmlckd** devices when the **devices** attribute is not set.

[Bugzilla:2187327](#)

8.11. COMPILERS AND DEVELOPMENT TOOLS

The **glibc system()** function now restores the previous signal mask unconditionally

Previously, if the **glibc system()** function was called concurrently from multiple threads, the signal mask for the **SIGCHLD** signal might not be restored correctly. As a consequence, the **SIGCHLD** signal remained blocked after the return from the **glibc system()** function on some threads.

With this update, the **glibc system()** function now restores the previous signal mask unconditionally, even when parallel **system()** function calls are running. As a result, the **SIGCHLD** signal is no longer incorrectly blocked if the **glibc system()** function is called concurrently from multiple threads.

[Bugzilla:2177235](#)

eu-addr2line -C now correctly recognizes other arguments

Previously, when you used the **-C** argument in **eu-addr2line** command from **elfutils**, the following single character argument disappeared. Consequently, the **eu-addr2line -Ci** command behaved the same way as **eu-addr2line -C** while **eu-addr2line -iC** worked as expected. This bug has been fixed, and **eu-addr2line -Ci** now recognizes both arguments.

[Bugzilla:2182059](#)

eu-addr2line -i now correctly handles code compiled with GCC link-time optimization

Previously, the **dwarf_getscopes** function from the **libdw** library included in **elfutils** was unable to find an abstract origin definition of a function that was compiled with GCC link-time optimization. Consequently, when you used the **-i** argument in the **eu-addr2line** command, **eu-addr2line** was unable to show inline functions for code compiled with **gcc -fllto**. With this update, the **libdw dwarf_getscopes** function looks in the correct compile unit for the inlined scope, and **eu-addr2line -i** works as expected.

[Bugzilla:2236182](#)

Programs using **papi** no longer stop when shutting down

Previously, **papi** initialized threads before **papi** initialized some components. Because of this, entries for certain components describing the number of elements in arrays were not set to correct values and zero-sized memory allocations were attempted. As a consequence, later accesses and frees of those zero-sized memory allocations caused the programs to stop.

The bug has been fixed and programs using **papi** no longer stop when shutting down.

[Bugzilla:2215582](#)

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:2186647](#)

8.12. IDENTITY MANAGEMENT

Paged searches from a regular user now do not impact performance

Previously, when Directory Server was under the search load, paged searches from a regular user could impact the server performance because a lock conflicted with the thread that polls for network events. In addition, if a network issue occurred while sending the page search, the whole server was unresponsive until the **nsslapd-iotimeout** parameter expired. With this update, the lock was split into several parts to avoid the contention with the network events. As a result, no performance impact during paged searches from a regular user.

[Bugzilla:1974242](#)

Schema replication now works correctly in Directory Server

Previously, when Directory Server replicated a schema to a new server, it added all the schema to the **99user.ldif** file on the remote replica. It seemed it was all custom schema because **X-ORIGIN** keyword was set to **user defined** for all definitions. As a result, it could cause issues with the web console and possibly for customers who monitor the schema and expect the **X-ORIGIN** keyword to have specific values. With this update, schema replication works as expected.

[Bugzilla:1759941](#)

Referral mode is now working correctly in Directory Server

Previously, CLI set **nsslapd-referral** configuration attribute to the backend and not to the mapping tree. As a result, referral mode did not work. With this update, the **nsslapd-referral** attribute is set correctly and the referral mode works as expected.

[Bugzilla:2053204](#)

The LMDB import now works faster

Previously, to build the **entryrdn** index, LMDB import worker threads waited for other worker threads to ensure that the parent entry was processed. This generated lock contention that drastically slowed import. With this update, the LDIF import over LMDB database was redesigned and the provider thread

stores the data about the entry RDN and its parents in a temporary database that the worker thread uses to build the **entryrdn** index. As a result, worker threads synchronization is no longer needed and the average import rate is better.

Note that the LMDB import still has an import rate three times slower than the BDB import because LMDB does not support concurrent write transactions.

[Bugzilla:2116948](#)

The **dirsrv** service now starts correctly after reboot

Previously, **dirsrv** service could fail to start after reboot because **dirsrv** service did not explicitly wait for **systemd-tmpfiles-setup.service** to finish. This led to a race condition. With this update, **dirsrv** service waits for the **systemd-tmpfiles-setup.service** to finish and no longer fail to start after reboot.

[Bugzilla:2179278](#)

Changing a security parameter now works correctly

Previously, when you changed a security parameter by using the **dsconf instance_name security set** command, the operation failed with the error:

```
Name 'log' is not defined
```

With this update, the security parameter change works as expected.

[Bugzilla:2189717](#)

SSSD now uses **sAMAccountName** when evaluating GPO-based access control

Previously, if **ldap_user_name** was set to a value other than **sAMAccountName** on an AD client, GPO-based access control failed. With this update, SSSD now always uses **sAMAccountName** when evaluating GPO-based access control. Even if **ldap_user_name** is set to a value different from **sAMAccountName** on an AD client, GPO-based access control now works correctly.

[Jira:SSSD-6107](#)

SSSD now handles duplicate attributes in the **user_attributes** option when retrieving users

Previously, if **sssd.conf** contained duplicate attributes in the **user_attributes** option, SSSD did not handle these duplicates correctly. As a consequence, users with those attributes could not be retrieved. With this update, SSSD now handles duplicates correctly. As a result, users with duplicate attributes can now be retrieved.

[Jira:SSSD-6177](#)

The dynamic Kerberos PAC ticket signature enforcement mechanism now fixes cross-version incompatibility in IdM

Previously, if your Identity Management (IdM) deployment featured servers running on both RHEL 9 and RHEL 8, the incompatibility caused by the upstream implementation of the Privilege Attribute Certificate (PAC) ticket signature support caused certain operations to fail. With this update, the implementation of the dynamic ticket signature enforcement mechanism feature in RHEL 9 fixes this cross-version incompatibility. For this feature to actually take effect, you must:

1. Update all the servers in the domain.
2. Restart all the IdM Kerberos Distribution Center (KDC) services.

The order of these two actions is important. When starting, the KDCs query the metadata of all the other servers in the domain to check if they all support the PAC ticket signature. If this is not the case, the signature will not be enforced.

For more information about the dynamic Kerberos PAC ticket signature enforcement mechanism, including an example of a constrained delegation request, see this [Knowledgebase article](#).

Jira:RHELDOCS-17011^[1], [Bugzilla:2182683](#), [Bugzilla:2178298](#)

SHA-1 signature verification can now be allowed in FIPS mode

Previously, it was not possible to allow the use of SHA-1 signature verification when Identity Management (IdM) was in FIPS mode. This is because IdM uses the FIPS-140-3 standard, which does not allow SHA-1 signatures. This situation caused problems with Active Directory (AD) interoperability, because AD only complies with the older FIPS-140-2 standard and therefore requires SHA-1 signatures.

This update introduces a FIPS exception for PKINIT signature verification. When FIPS mode is enabled in IdM, its restrictions are ignored. Only default mode restrictions are applied, allowing the use of the **SHA1** crypto module even when in FIPS mode. As a result, AD interoperability in FIPS mode works as intended.

In the scenario of an IdM/AD trust, or using a RHEL 9.2 or later host as an AD client, you need to set the crypto policy to FIPS:AD-SUPPORT:SHA1 to support PKINIT while in FIPS mode.

[Bugzilla:2155607](#)

Deleting the IdM admin user is now no longer permitted

Previously, nothing prevented you from deleting the Identity Management (IdM) **admin** user if you were a member of the **admins** group. The absence of the **admin** user causes the trust between IdM and Active Directory (AD) to stop functioning correctly. With this update, you can no longer delete the **admin** user. As a result, the IdM-AD trust works correctly.

[Bugzilla:2229712](#)

ipa-kdb no longer causes krb5kdc to fail

Previously, the **ipa-kdb** driver did not differentiate between the absence of a server host object and a connection failure. Consequently, the **krb5kdc** server sometimes stopped unexpectedly because of a **NULL** LDAP context produced by a connection issue with the LDAP server.

With this update, the **ipa-kdb** driver correctly identifies connection failures and differentiates between them and the absence of a server host object. As a result, the **krb5kdc** server does not fail anymore.

[Bugzilla:2227831](#)

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 truststore and the IdM client installer does not set up the TLS CA configuration in the **ldap.conf** file.

[Bugzilla:2094673](#)

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

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

With the release of [RHBA-2023:4359](#), a case-sensitive comparison is replaced with a case-insensitive comparison that ignores the case of a character. As a result, IdM clients can now lookup users of an AD trusted domain, even if their usernames contain mixed case characters and they are configured with overrides in IdM.

Jira:SSSD-6096

8.13. THE WEB CONSOLE

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

In RHEL 9.2, 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 this update, 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:2203361](#)

VNC console now works at most resolutions

Previously, when using the Virtual Network Computing (VNC) console under certain display resolutions, a mouse offset problem was present or only a part of the interface was visible. Consequently, using the VNC console was not possible.

With this update, the problem has been fixed and the VNC console works correctly at most resolutions, with the exception of ultra high resolutions, such as 3840x2160.

Note that a small offset between the recorded and displayed positions of the cursor might still be present. However, this does not significantly impact the usability of the VNC console.

[Bugzilla:2030836](#)

8.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The storage role can now resize the mounted file systems without unmounting

Previously, the **storage** role was unable to resize mounted devices, even if the file system supported online resizing. As a consequence, the **storage** role unmounted all file systems before resizing, which failed for file systems that were in use, for example, while resizing the `/` directory of the running system.

With this update, the **storage** role now supports resizing mounted file systems that support online resizing such as XFS and Ext4. As a result, the mounted file systems can now be resized without unmounting them.

[Bugzilla:2168692](#)

The `podman_registries_conf` variable now configures `unqualified-search-registries` field correctly

Previously, after configuring the `podman_registries_conf` variable, the `podman` RHEL system role failed. Consequently, `unqualified-search-registries = ["registry.access.redhat.com"]` setting was not generated in the `/etc/containers/registries.conf.d/50-systemroles.conf` file. With this update, this problem has been fixed.

[Bugzilla:2211984](#)

The `kdump` role adds `authorized_keys` idempotently

Previously, the task to add `authorized_key` added an extra newline character every time. Consequently the role was not acting idempotent. With this fix, adding a new `authorized_key` works correctly and adds only a single key value idempotently.

[Bugzilla:2232241](#)

The `kdump` system role does not fail if `kdump_authorized_keys` is missing

Previously, the `kdump` system role failed to add `SSH` authorized keys if the user defined in the `kdump_ssh_user` variable did not have access to the `.ssh` directory in the `home` directory or an empty `.ssh/authorized_keys` file. With this fix, the `kdump` system role now correctly adds authorized keys to the `SSH` configuration. As a result, the key based authentication works reliably in the described scenario.

[Bugzilla:2232231](#)

Failure to remove data from member disks before creation no longer persists

Previously, when creating RAID volumes, the system did not effectively eliminate existing data from member disks before forming the RAID volume. With this update, RAID volumes remove any pre-existing data from member disks as needed.

[Bugzilla:2224090](#)

Running the `firewall` RHEL system role in check mode with non-existent services no longer fails

Previously, running the `firewall` role in check mode with non-existent services would fail. This fix implements better compliance with Ansible best practices for check mode. As a result, non-existent services being enabled or disabled no longer fails the role in check mode. Instead, a warning prompts you to confirm that the service is defined in a previous playbook.

[Bugzilla:2222428](#)

The `firewall` RHEL system role on RHEL 7 no longer attempts to install non-existent Python packages

Previously, when the `firewall` role on RHEL 7 was called from another role, and that role was using `python3`, the `firewall` role attempted to install the `python3-firewall` library for that version of Python. However, that library is not available in RHEL 7. Consequently, the `python3-firewall` library was not found, and you received the following error message:

```
No package matching 'python3-firewall' found available, installed or updated
```

With this update, the **firewall** role does not attempt to install the **python-firewall** or **python3-firewall** library. As a result, the **firewall** role does not fail on RHEL 7 when **python3** is installed on the managed node.

[Bugzilla:2216520](#)

kdump RHEL system role updates

The **kdump** RHEL system role has been updated to a newer version, which brings the following notable enhancements:

- After installing **kexec-tools**, the utility suite no longer generates the **/etc/sysconfig/kdump** file because you do not need to manage this file anymore.
- The role supports the **auto_reset_crashkernel** and **dracut_args** variables.

For more details, see resources in the **/usr/share/doc/rhel-system-roles/kdump/** directory.

[Bugzilla:2211187](#)

Insights tags created by using the **rhc role are now applied correctly**

Previously, when you created Insights tags by using the **rhc** role, tags were not stored in the correct file. Consequently, tags were not sent to Insights and as a result they were not applied to the systems in the Insights inventory.

With this fix, tags are stored correctly and applied to the systems present in the Insights inventory.

[Bugzilla:2209200](#)

raid_chunk_size parameter no longer returns an error message

Previously, **raid_chunk_size** attribute was not allowed for RAID pools and volumes. With this update, you can now configure the **raid_chunk_size** attribute for RAID pools and volumes without encountering any restrictions.

[Bugzilla:2193058](#)

The **certificate RHEL system role now checks for the certificate key size when determining whether to perform a new certificate request**

Previously, the **certificate** RHEL system role did not check the key size of a certificate when evaluating whether to request a new certificate. As a consequence, the role sometimes did not issue new certificate requests in cases where it should. With this update, **certificate** now checks the **key_size** parameter to determine if a new certificate request should be performed.

[Bugzilla:2186057](#)

The **kdump role adds multiple keys to **authorized_keys** idempotently**

Previously, adding multiple SSH keys to the **authorized_keys** file at the same time replaced the key value of one host by another. This update fixes the problem by using the **lineinfile** module to manage the **authorized_keys** file. **lineinfile** iterates the tasks in sequence, checking for an existing key and writing the new key in one atomic operation on a single host at one time. As a result, adding SSH keys on multiple hosts works correctly, and does not replace the key value from another host.

Note: Use the **serial: 1** play serial keyword at play level to control the number of hosts executing at one time.

Jira:RHEL-1499^[1]

The **kdump** role successfully updates `.ssh/authorized_keys` for **kdump_ssh_server** authentication

Previously, the `.ssh` directory was not accessible by the **kdump** role to securely authenticate users to log into **kdump_ssh_server**. As a consequence, the **kdump** role did not update the `.ssh/authorized_keys` file and the SSH mechanism to verify the **kdump_ssh_server** failed. This update fixes the problem. As a result the **kdump_ssh_user** authentication on **kdump_ssh_server** works reliably.

Jira:RHEL-1397^[1]

Enabling **kdump** for system role requires using the `failure_action` configuration parameter on RHEL 9 and later versions

Previously, using the **default** option during **kdump** configuration was not successful and printed the following warning in logs:

```
kdump: warning: option 'default' was renamed 'failure_action' and will be removed in the future.  
please update /etc/kdump.conf to use option 'failure_action' instead.
```

Consequently, the role did not enable **kdump** successfully if **default** option was used. This update fixes the problem and you can configure kernel dump parameters on multiple systems by using the **failure_action** parameter. As a result, enabling **kdump** works successfully in the described scenario.

Jira:RHEL-906^[1]

The **previous: replaced** parameter of the **firewall** system role now overrides the previous configuration without deleting it

Previously, if you added the **previous: replaced** parameter to the variable list, the **firewall** system role removed all existing user-defined settings and reset **firewalld** to the default settings. This fix uses the fallback configuration in **firewalld**, which was introduced in the EL7 release, to retain the previous configuration. As a result, when you use the **previous: replaced** parameter in the variable list, the **firewall.conf** configuration file is not deleted on reset, but the file and comments in the file are retained.

Jira:RHEL-1495^[1]

The **firewall** RHEL system role correctly reports changes when using **previous: replaced** in check mode

Previously, the **firewall** role was not checking whether any files would be changed when using the **previous: replaced** parameter in check mode. As a consequence, the role gave an error about undefined variables. This fix adds new check variables to the check mode to assess whether any files would be changed by the **previous: replaced** parameter. The check for the **firewalld.conf** file assesses the **rpm** database to determine whether the file has been changed from the version shipped in the package. As a result, the **firewall** role now correctly reports changes when using the **previous: replaced** parameter.

Jira:RHEL-898^[1]

The **firewall** RHEL system role correctly reports changes when assigning zones to Network Manager interfaces

Previously, the Network Manager interface assignment reported changes when no changes were present. With this fix, the **try_set_zone_of_interface** module in the file **library/firewall_lib.py** returns a

second value, which denotes whether the interface's zone was changed. As a result, the module now correctly reports changes when assigning zones to interfaces handled by Network Manager.

[Jira:RHEL-885^{\[1\]}](#)

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

The NVIDIA graphics device continues working after VM shutdown

Previously, in the RHEL kernel, device power transition delays were more closely aligned to those required by the PCIe specification. As a consequence, some NVIDIA GPUs could become unresponsive when used for device assignment after a shutdown of the attached VM. This update extends the device power transition delay for NVIDIA audio device functions. As a result, NVIDIA GPUs continue to work correctly in this scenario.

[Bugzilla:2178956^{\[1\]}](#)

Failover virtio NICs are now correctly assigned an IP address on Windows virtual machines

Previously, when starting a Windows virtual machine (VM) with only a failover virtio NIC, the VM failed to assign an IP address to the NIC. Consequently, the NIC was unable to set up a network connection. This problem has been fixed and VM NICs now set up network connections as expected in the described scenario.

[Bugzilla:1969724](#)

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^{\[1\]}](#)

Broadcom network adapters now work correctly on Windows VMs after a live migration

Previously, network adapters from the Broadcom family of devices, such as Broadcom, Qlogic, or Marvell, could not be hot-unplugged during live migration of Windows virtual machines (VMs). As a consequence, the adapters worked incorrectly after the migration was complete. This problem affected only adapters that were attached to Windows VMs using Single-root I/O virtualization (SR-IOV). With this update, the underlying code has been fixed and the problem no longer occurs.

[Jira:RHEL-910](#), [Bugzilla:2091528](#), [Bugzilla:2111319](#)

nodedev-dumpxml lists attributes correctly for certain mediated devices

Before this update, the **nodedev-dumpxml** utility did not list attributes correctly for mediated devices that were created using the **nodedev-create** command. This has been fixed, and **nodedev-dumpxml** now displays the attributes of the affected mediated devices properly.

[Bugzilla:2143158](#)

virtiofs devices could not be attached after restarting virtqemud or libvirtd

Previously, restarting the **virtqemud** or **libvirtd** services prevented **virtiofs** storage devices from being attached to virtual machines (VMs) on your host. This bug has been fixed, and you can now attach **virtiofs** devices in the described scenario as expected.

[Bugzilla:2078693](#)

Hot plugging a Watchdog card to a virtual machine no longer fails

Previously, if no PCI slots were available, adding a Watchdog card to a running virtual machine (VM) failed with the following error:

```
Failed to configure watchdog
ERROR Error attempting device hotplug: internal error: No more available PCI slots
```

With this update, the problem has been fixed and adding a Watchdog card to a running VM now works as expected.

[Bugzilla:2173584](#)

blob resources do not work correctly for virtio-gpu on IBM Z

The **virtio-gpu** device is currently not compatible with **blob** memory resources on IBM Z systems. As a consequence, if you configure a virtual machine (VM) with **virtio-gpu** on an IBM Z host to use **blob** resources, the VM does not have any graphical output.

[Jira:RHEL-7135](#)

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 installation program as a Technology Preview

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

[Bugzilla:2107346](#)

9.2. 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:2108532^{\[1\]}](#)

9.3. 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^{\[1\]}](#)

9.4. INFRASTRUCTURE SERVICES

Socket API for Tuned available as a Technology Preview

The socket API for controlling Tuned through a 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](#)

9.5. 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](#)^[1]

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](#)^[1]

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](#)

The PRP and HSR protocols are now available as a Technology Preview

This update adds the **hsr** kernel module that provides the following protocols:

- Parallel Redundancy Protocol (PRP)
- High-availability Seamless Redundancy (HSR)

The IEC 62439-3 standard defines these protocols, and you can use this feature to configure zero-loss redundancy in Ethernet networks.

[Bugzilla:2177256](#)^[1]

Offloading IPsec encapsulation to a NIC is now available as a Technology Preview

This update adds the IPsec packet offloading capabilities to the kernel. Previously, it was possible to only offload the encryption to a network interface controller (NIC). With this enhancement, the kernel can now offload the entire IPsec encapsulation process to a NIC to reduce the workload.

Note that offloading the IPsec encapsulation process to a NIC also reduces the ability of the kernel to monitor and filter such packets.

Bugzilla:2178699^[1]

Network drivers for modems in RHEL are available as Technology Preview

Device manufacturers support Federal Communications Commission (FCC) locking as the default setting. FCC provides a lock to bind WWAN drivers to a specific system where WWAN drivers provide a channel to communicate with modems. Based on the modem PCI ID, manufacturers integrate unlocking tools on Red Hat Enterprise Linux for ModemManager. However, a modem remains unusable if not unlocked previously even if the WWAN driver is compatible and functional. Red Hat Enterprise Linux provides the drivers for the following modems with limited functionality as a Technology Preview:

- Qualcomm MHI WWAN MBIM - Telit FN990Axx
- Intel IPC over Shared Memory (IOSM) - Intel XMM 7360 LTE Advanced
- Mediatek t7xx (WWAN) - Fibocom FM350GL
- Intel IPC over Shared Memory (IOSM) - Fibocom L860GL modem

Jira:RHELDPCS-16760^[1], Bugzilla:2123542, Jira:RHEL-6564, Bugzilla:2110561, Bugzilla:2222914

Segment Routing over IPv6 (SRv6) is available as a Technology Preview

The RHEL kernel provides Segment Routing over IPv6 (SRv6) as a Technology Preview. You can use this functionality to optimize traffic flows in edge computing or to improve network programmability in data centers. However, the most significant use case is the end-to-end (E2E) network slicing in 5G deployment scenarios. In that area, the SRv6 protocol provides you with the programmable custom network slices and resource reservations to address network requirements for specific applications or services. At the same time, the solution can be deployed on a single-purpose appliance, and it satisfies the need for a smaller computational footprint.

Bugzilla:2186375^[1]

kTLS rebased to version 6.3

The kernel Transport Layer Security (kTLS) functionality is a Technology Preview. With this RHEL release, kTLS has been rebased to the 6.3 upstream version, and notable changes include:

- Added the support for 256-bit keys with TX device offload
- Delivered various bugfixes

Bugzilla:2183538^[1]

9.6. KERNEL

The **kdump** mechanism with a unified kernel image is available as a Technology Preview

The **kdump** mechanism with a kernel image contained in a unified kernel image (UKI) is available as a Technology Preview. UKI is a single executable, combining the **initramfs**, **vmlinuz**, and the kernel command line in a single file. The UKI key benefit being extending the cryptographic signature for SecureBoot to all components at once.

For the feature to work, with the kernel command line contained in the UKI, set the **crashkernel=** parameter with an appropriate value. This reserves the required memory for **kdump**.

Note: Currently the **kexec_file_load** system call from the Linux kernel cannot load UKI. Therefore, only the kernel image contained in the UKI is used when loading the crash kernel with the **kexec_file_load** system call.

[Bugzilla:2169720^{\[1\]}](#)

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:1874182^{\[1\]}](#)

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^{\[1\]}](#)

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^[1]

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^[1]

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^[1]

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^[1]

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^[1]

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^[1]

The `io_uring` interface is available as a Technology Preview

io_uring is a new and effective asynchronous I/O interface, which is now available as a Technology Preview. By default, this feature is disabled. You can enable this interface by setting the **kernel.io_uring_disabled** sysctl variable to any one of the following values:

0

All processes can create **io_uring** instances as usual.

1

io_uring creation is disabled for unprivileged processes. The **io_uring_setup** fails with the **-EPERM** error unless the calling process is privileged by the **CAP_SYS_ADMIN** capability. Existing **io_uring** instances can still be used.

2

io_uring creation is disabled for all processes. The **io_uring_setup** always fails with **-EPERM**. Existing **io_uring** instances can still be used. This is the default setting.

An updated version of the SELinux policy to enable the **mmap** system call on anonymous inodes is also required to use this feature.

By using the **io_uring** command pass-through, an application can issue commands directly to the underlying hardware, such as **nvme**. Use of **io_uring** command pass-through currently requires a custom SELinux policy module. Create a custom SELinux policy module:

1. Save the following lines as **io_uring_cmd_passthrough.cil** file:

```
---cut here---
( allow unconfined_domain_type device_node ( io_uring ( cmd )))
( allow unconfined_domain_type file_type ( io_uring ( cmd )))
---cut here---
```

2. Load the policy module:

```
# semodule -i io_uring_cmd_passthrough.cil
```

Bugzilla:2068237^[1]

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, and 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](#)

ACME available as a Technology Preview

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

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



IMPORTANT

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



WARNING

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

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

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

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

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

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

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

Bugzilla:2084181^[1]

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^[1]

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^[1]

9.11. VIRTUALIZATION

Creating nested virtual machines

Nested KVM virtualization is provided as a Technology Preview for KVM virtual machines (VMs) running on Intel, AMD64, and IBM Z hosts with RHEL 9. With this feature, a RHEL 7, RHEL 8, or RHEL 9 VM that runs on a physical RHEL 9 host can act as a hypervisor, and host its own VMs.

Jira:RHELDPCS-17040^[1]

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^[1]

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^[1]

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.

Note, however, that **virtio-mem** currently does not work on VMs that use a Windows operating system.

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

Intel TDX in RHEL guests

As a Technology Preview, the Intel Trust Domain Extension (TDX) feature can now be used in RHEL 9.2 and later 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](#)^[1]

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](#)^[1]

Intel vGPU available as a Technology Preview

As a Technology Preview, it is possible to divide a physical Intel GPU device into multiple virtual devices referred to as **mediated devices**. These mediated devices can then be assigned to multiple virtual machines (VMs) as virtual GPUs. As a result, these VMs share the performance of a single physical Intel GPU.

Note that this feature is deprecated and was removed entirely with the RHEL 9.3 release.

[Jira:RHELDPCS-17050](#)^[1]

9.12. 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](#)^[1]

9.13. CONTAINERS

SQLite database backend for Podman is available as a Technology Preview

Beginning with Podman v4.6, the SQLite database backend for Podman is available as a Technology Preview. To set the database backend to SQLite, add the **database_backend = "sqlite"** option in the **/etc/containers/containers.conf** configuration file. Run the **podman system reset** command to reset storage back to the initial state before you switch to the SQLite database backend. Note that you have to re-create all containers and pods. The SQLite database guarantees good stability and consistency. Other databases in the containers stack will be moved to SQLite as well. The BoltDB remains the default database backend.

Jira:RHELPLAN-154429^[1]

The **podman-machine** command is unsupported

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

Jira:RHELDPCS-16861^[1]

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.

Deprecated devices are fully supported, that means that they are tested and maintained, and their support status 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. 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 --nontp**
- **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^[1]

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.

Note that specifying a user or group customization in blueprints that are used to deploy an existing OSTree commit, such as **edge-raw-image**, **edge-installer**, and **edge-simplified-installer** image types remains supported.

[Bugzilla:2173928](#)

The **initial-setup** package now has been deprecated

The **initial-setup** package has been deprecated in Red Hat Enterprise Linux 9.3 and will be removed in the next major RHEL release. As a replacement, use **gnome-initial-setup** for the graphical user interface.

Jira:RHELDPCS-16393^[1]

The **provider_hostip** and **provider_fedora_geoip** values of the **inst.geoloc** boot option are deprecated

The **provider_hostip** and **provider_fedora_geoip** values that specified the GeoIP API for the **inst.geoloc=** boot option are deprecated. As a replacement, you can use the **geolocation_provider=URL** option to set the required geolocation in the installation program configuration file. You can still use the **inst.geoloc=0** option to disable the geolocation.

[Bugzilla:2127473](#)

10.2. 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^[1]

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^[1]

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](#)

NTLM and Krb4 are deprecated in Cyrus SASL

The NTLM and Kerberos 4 authentication protocols have been deprecated and might be removed in a future major version of RHEL. These protocols are no longer considered secure and have already been removed from upstream implementations.

Jira:RHELDOCS-17380^[1]

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^{\[1\]}](#)

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^[1]

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

10.3. 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

10.4. SHELLS AND COMMAND-LINE TOOLS

Setting the `TMPDIR` variable in the ReaR configuration file is deprecated

Setting the `TMPDIR` environment variable in the `/etc/rear/local.conf` or `/etc/rear/site.conf` ReaR configuration file), by using a statement such as `export TMPDIR=...`, does not work and is deprecated.

To specify a custom directory for ReaR temporary files, export the variable in the shell environment before executing ReaR. For example, execute the `export TMPDIR=...` statement and then execute the `rear` command in the same shell session or script.

Jira:RHELDPCS-18049

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^[1]

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.

[Jira:RHEL-6856](#)

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^{\[1\]}](#)

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^{\[1\]}](#)

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](#)

The PF_KEYv2 kernel API is deprecated

Applications can configure the kernel's IPsec implementation by using the **PV_KEYv2** and the newer **netlink** API. **PV_KEYv2** is not actively maintained upstream and misses important security features, such as modern ciphers, offload, and extended sequence number support. As a result, starting with RHEL 9.3, the **PV_KEYv2** API is deprecated and will be removed in the next major RHEL release. If you use this kernel API in your application, migrate it to use the modern **netlink** API as an alternative.

[Jira:RHEL-1015^{\[1\]}](#)

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.

For more information, see [Is kexec_load supported in RHEL9?](#) .

[Bugzilla:2113873^{\[1\]}](#)

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^{\[1\]}](#)

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](#)

Persistent Memory Development Kit (`pmdk`) and support library have been deprecated in RHEL 9

`pmdk` is a collection of libraries and tools for System Administrators and Application Developers to simplify managing and accessing persistent memory devices. `pmdk` and support library have been deprecated in RHEL 9. This also includes the `-debuginfo` packages.

The following list of binary packages produced by **pmdk**, including the **nvml** source package have been deprecated:

- **libpmem**
- **libpmem-devel**
- **libpmem-debug**
- **libpmem2**
- **libpmem2-devel**
- **libpmem2-debug**
- **libpmemblk**
- **libpmemblk-devel**
- **libpmemblk-debug**
- **libpmemlog**
- **libpmemlog-devel**
- **libpmemlog-debug**
- **libpmemobj**
- **libpmemobj-devel**
- **libpmemobj-debug**
- **libpmempool**
- **libpmempool-devel**
- **libpmempool-debug**
- **pmempool**
- **daxio**
- **pmreorder**
- **pmdk-convert**
- **libpmemobj++**
- **libpmemobj++-devel**
- **libpmemobj++-doc**

Jira:RHELDOS-16432^[1]

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^[1], Jira:RHELPLAN-80695, [Bugzilla:1974657](#)

10.9. COMPILERS AND DEVELOPMENT TOOLS

Smaller size of keys than 2048 are deprecated by openssl 3.0 in Go's FIPS mode

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 in Go's FIPS mode

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^{\[1\]}](#)

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^{\[1\]}](#)

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^{\[1\]}](#)

The nsslapd-ldapimaprootdn parameter is deprecated

In Directory Server, the **nsslapd-ldapimaprootdn** configuration parameter is used to map a system root entry to a root DN entry. Usually, the **nsslapd-ldapimaprootdn** parameter has the same value as the **nsslapd-rootdn** parameter. In addition, changing one attribute but not changing the other leads to a non-functional autobind configuration that breaks **dsconf** utility and access to the web console.

With this update, Directory Server uses only the **nsslapd-rootdn** parameter to map a system root entry to a root DN entry. As a result, the **nsslapd-ldapimaprootdn** parameter is deprecated and the root DN change does not break **dsconf** utility and access to the web console.

[Bugzilla:2170494](#)

The nsslapd-conntablesizesize configuration parameter has been removed from 389-ds-base

The **nsslapd-conntablesizesize** configuration parameter has been removed from the **389-ds-base** package in RHEL 9.3. Previously, the **nsslapd-conntablesizesize** configuration attribute specified the size of the connection table that managed established connections. With the introduction of the multi-listener feature, which improves the management of established connections, Directory Server now calculates the size of the connection table dynamically. This also resolves issues, when the connection table size was set too low and it affected the number of connections the server was able to support. Starting with RHEL 9.3, use only **nsslapd-maxdescriptors** and **nsslapd-reservedescriptors** attributes to manage the number of TCP/IP connections Directory Server can support.

[Bugzilla:2098236](#)

The SMB1 protocol is deprecated in Samba

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

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

Jira:RHELDOCS-16612^[1]

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^[1]

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:RHELDOCS-16300^[1]

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^[1]

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^[1]

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 might 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:RHELDPCS-16948^[1], 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** might not be yet available in the RHEL web console.

Jira:RHELPLAN-10304^[1]

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^[1]

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 before Intel Xeon 55xx and 75xx Processor families (also known as Nehalem)
- For AMD: models before AMD Opteron G4
- For IBM Z: models before 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^[1]

The Intel vGPU feature has been removed

Previously, as a Technology Preview, it was possible to divide a physical Intel GPU device into multiple virtual devices referred to as **mediated devices**. These mediated devices could then be assigned to multiple virtual machines (VMs) as virtual GPUs. As a result, these VMs shared the performance of a single physical Intel GPU, however only selected Intel GPUs were compatible with this feature.

Since RHEL 9.3, the Intel vGPU feature has been removed entirely.

[Bugzilla:2206599^{\[1\]}](#)

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^{\[1\]}](#)

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^{\[1\]}](#)

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 is deprecated and will be removed from Podman in a future minor release of RHEL. 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:RHELDOCS-16756^{\[1\]}](#)

The Inkscape and LibreOffice Flatpak images are deprecated

The **rhel9/inkscape-flatpak** and **rhel9/libreoffice-flatpak** Flatpak images, which are available as Technology Previews, have been deprecated.

Red Hat recommends the following alternatives to these images:


- To replace **rhel9/inkscape-flatpak**, use the **inkscape** RPM package.
- To replace **rhel9/libreoffice-flatpak**, see the [LibreOffice deprecation release note](#) .

[Jira:RHELDOCS-17102^{\[1\]}](#)

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:

- adwaita-gtk2-theme
- autocorr-af
- autocorr-bg
- autocorr-ca
- autocorr-cs
- autocorr-da
- autocorr-de
- autocorr-dsb
- autocorr-el
- autocorr-en
- autocorr-es
- autocorr-fa
- autocorr-fi
- autocorr-fr
- autocorr-ga
- autocorr-hr
- autocorr-hsb
- autocorr-hu
- autocorr-is
- autocorr-it
- autocorr-ja
- autocorr-ko

- autocorr-lb
- autocorr-lt
- autocorr-mn
- autocorr-nl
- autocorr-pl
- autocorr-pt
- autocorr-ro
- autocorr-ru
- autocorr-sk
- autocorr-sl
- autocorr-sr
- autocorr-sv
- autocorr-tr
- autocorr-vi
- autocorr-vro
- autocorr-zh
- cheese
- cheese-libs
- clutter
- clutter-gst3
- clutter-gtk
- cogl
- daxio
- dbus-glib
- dbus-glib-devel
- enchant
- enchant-devel
- eog
- evolution

- evolution-bogofilter
- evolution-devel
- evolution-help
- evolution-langpacks
- evolution-mapi
- evolution-mapi-langpacks
- evolution-pst
- evolution-spamassassin
- festival
- festival-data
- festvox-slt-arctic-hts
- flite
- flite-devel
- gedit
- gedit-plugin-bookmarks
- gedit-plugin-bracketcompletion
- gedit-plugin-codecomment
- gedit-plugin-colorpicker
- gedit-plugin-colorschemer
- gedit-plugin-commander
- gedit-plugin-drawspaces
- gedit-plugin-findinfiles
- gedit-plugin-joinlines
- gedit-plugin-multiedit
- gedit-plugin-sessionsaver
- gedit-plugin-smartspaces
- gedit-plugin-synctex
- gedit-plugin-terminal
- gedit-plugin-textsize

- gedit-plugin-translate
- gedit-plugin-wordcompletion
- gedit-plugins
- gedit-plugins-data
- gnome-common
- gnome-photos
- gnome-photos-tests
- gnome-screenshot
- gnome-themes-extra
- gtk2
- gtk2-devel
- gtk2-devel-docs
- gtk2-immodule-xim
- gtk2-immodules
- highcontrast-icon-theme
- inkscape
- inkscape-docs
- inkscape-view
- iptables-devel
- iptables-libs
- iptables-nft
- iptables-nft-services
- iptables-utils
- libdb
- libgdata
- libgdata-devel
- libpmem
- libpmem-debug
- libpmem-devel

- libpmem2
- libpmem2-debug
- libpmem2-devel
- libpmemblk
- libpmemblk-debug
- libpmemblk-devel
- libpmemlog
- libpmemlog-debug
- libpmemlog-devel
- libpmemobj
- libpmemobj-debug
- libpmemobj-devel
- libpmempool
- libpmempool-debug
- libpmempool-devel
- libreoffice
- libreoffice-base
- libreoffice-calc
- libreoffice-core
- libreoffice-data
- libreoffice-draw
- libreoffice-emailmerge
- libreoffice-filters
- libreoffice-gdb-debug-support
- libreoffice-graphicfilter
- libreoffice-gtk3
- libreoffice-help-ar
- libreoffice-help-bg
- libreoffice-help-bn

- libreoffice-help-ca
- libreoffice-help-cs
- libreoffice-help-da
- libreoffice-help-de
- libreoffice-help-dz
- libreoffice-help-el
- libreoffice-help-en
- libreoffice-help-eo
- libreoffice-help-es
- libreoffice-help-et
- libreoffice-help-eu
- libreoffice-help-fi
- libreoffice-help-fr
- libreoffice-help-gl
- libreoffice-help-gu
- libreoffice-help-he
- libreoffice-help-hi
- libreoffice-help-hr
- libreoffice-help-hu
- libreoffice-help-id
- libreoffice-help-it
- libreoffice-help-ja
- libreoffice-help-ko
- libreoffice-help-lt
- libreoffice-help-lv
- libreoffice-help-nb
- libreoffice-help-nl
- libreoffice-help-nn
- libreoffice-help-pl

- libreoffice-help-pt-BR
- libreoffice-help-pt-PT
- libreoffice-help-ro
- libreoffice-help-ru
- libreoffice-help-si
- libreoffice-help-sk
- libreoffice-help-sl
- libreoffice-help-sv
- libreoffice-help-ta
- libreoffice-help-tr
- libreoffice-help-uk
- libreoffice-help-zh-Hans
- libreoffice-help-zh-Hant
- libreoffice-impress
- libreoffice-langpack-af
- libreoffice-langpack-ar
- libreoffice-langpack-as
- libreoffice-langpack-bg
- libreoffice-langpack-bn
- libreoffice-langpack-br
- libreoffice-langpack-ca
- libreoffice-langpack-cs
- libreoffice-langpack-cy
- libreoffice-langpack-da
- libreoffice-langpack-de
- libreoffice-langpack-dz
- libreoffice-langpack-el
- libreoffice-langpack-en
- libreoffice-langpack-eo

- libreoffice-langpack-es
- libreoffice-langpack-et
- libreoffice-langpack-eu
- libreoffice-langpack-fa
- libreoffice-langpack-fi
- libreoffice-langpack-fr
- libreoffice-langpack-fy
- libreoffice-langpack-ga
- libreoffice-langpack-gl
- libreoffice-langpack-gu
- libreoffice-langpack-he
- libreoffice-langpack-hi
- libreoffice-langpack-hr
- libreoffice-langpack-hu
- libreoffice-langpack-id
- libreoffice-langpack-it
- libreoffice-langpack-ja
- libreoffice-langpack-kk
- libreoffice-langpack-ko
- libreoffice-langpack-lt
- libreoffice-langpack-lv
- libreoffice-langpack-mai
- libreoffice-langpack-ml
- libreoffice-langpack-mr
- libreoffice-langpack-nb
- libreoffice-langpack-nl
- libreoffice-langpack-nn
- libreoffice-langpack-nr

- libreoffice-langpack-nso
- libreoffice-langpack-or
- libreoffice-langpack-pa
- libreoffice-langpack-pl
- libreoffice-langpack-pt-BR
- libreoffice-langpack-pt-PT
- libreoffice-langpack-ro
- libreoffice-langpack-ru
- libreoffice-langpack-si
- libreoffice-langpack-sk
- libreoffice-langpack-sl
- libreoffice-langpack-sr
- libreoffice-langpack-ss
- libreoffice-langpack-st
- libreoffice-langpack-sv
- libreoffice-langpack-ta
- libreoffice-langpack-te
- libreoffice-langpack-th
- libreoffice-langpack-tn
- libreoffice-langpack-tr
- libreoffice-langpack-ts
- libreoffice-langpack-uk
- libreoffice-langpack-ve
- libreoffice-langpack-xh
- libreoffice-langpack-zh-Hans
- libreoffice-langpack-zh-Hant
- libreoffice-langpack-zu
- libreoffice-math
- libreoffice-ogltrans

- libreoffice-opensymbol-fonts
- libreoffice-pdfimport
- libreoffice-pyuno
- libreoffice-sdk
- libreoffice-sdk-doc
- libreoffice-ure
- libreoffice-ure-common
- libreoffice-wiki-publisher
- libreoffice-writer
- libreoffice-x11
- libreoffice-xsltfilter
- libreofficekit
- libsoup
- libsoup-devel
- libuser
- libuser-devel
- libwpe
- libwpe-devel
- mcpp
- mod_auth_mellon
- motif
- motif-devel
- pmdk-convert
- pmempool
- python3-pytz
- qt5
- qt5-assistant
- qt5-designer
- qt5-devel

- qt5-doctools
- qt5-linguist
- qt5-qdbusviewer
- qt5-qt3d
- qt5-qt3d-devel
- qt5-qt3d-doc
- qt5-qt3d-examples
- qt5-qtbase
- qt5-qtbase-common
- qt5-qtbase-devel
- qt5-qtbase-doc
- qt5-qtbase-examples
- qt5-qtbase-gui
- qt5-qtbase-mysql
- qt5-qtbase-odbc
- qt5-qtbase-postgresql
- qt5-qtbase-private-devel
- qt5-qtbase-static
- qt5-qtconnectivity
- qt5-qtconnectivity-devel
- qt5-qtconnectivity-doc
- qt5-qtconnectivity-examples
- qt5-qtdeclarative
- qt5-qtdeclarative-devel
- qt5-qtdeclarative-doc
- qt5-qtdeclarative-examples
- qt5-qtdeclarative-static
- qt5-qtdoc
- qt5-qtgraphicaleffects

- [qt5-qtgraphicaleffects-doc](#)
- [qt5-qtimageformats](#)
- [qt5-qtimageformats-doc](#)
- [qt5-qtlocation](#)
- [qt5-qtlocation-devel](#)
- [qt5-qtlocation-doc](#)
- [qt5-qtlocation-examples](#)
- [qt5-qtmultimedia](#)
- [qt5-qtmultimedia-devel](#)
- [qt5-qtmultimedia-doc](#)
- [qt5-qtmultimedia-examples](#)
- [qt5-qtquickcontrols](#)
- [qt5-qtquickcontrols-doc](#)
- [qt5-qtquickcontrols-examples](#)
- [qt5-qtquickcontrols2](#)
- [qt5-qtquickcontrols2-devel](#)
- [qt5-qtquickcontrols2-doc](#)
- [qt5-qtquickcontrols2-examples](#)
- [qt5-qtscript](#)
- [qt5-qtscript-devel](#)
- [qt5-qtscript-doc](#)
- [qt5-qtscript-examples](#)
- [qt5-qtsensors](#)
- [qt5-qtsensors-devel](#)
- [qt5-qtsensors-doc](#)
- [qt5-qtsensors-examples](#)
- [qt5-qtserialbus](#)
- [qt5-qtserialbus-devel](#)
- [qt5-qtserialbus-doc](#)

- qt5-qtserialbus-examples
- qt5-qtserialport
- qt5-qtserialport-devel
- qt5-qtserialport-doc
- qt5-qtserialport-examples
- qt5-qtsvg
- qt5-qtsvg-devel
- qt5-qtsvg-doc
- qt5-qtsvg-examples
- qt5-qttools
- qt5-qttools-common
- qt5-qttools-devel
- qt5-qttools-doc
- qt5-qttools-examples
- qt5-qttools-libs-designer
- qt5-qttools-libs-designercomponents
- qt5-qttools-libs-help
- qt5-qttools-static
- qt5-qttranslations
- qt5-qtwayland
- qt5-qtwayland-devel
- qt5-qtwayland-doc
- qt5-qtwayland-examples
- qt5-qtwebchannel
- qt5-qtwebchannel-devel
- qt5-qtwebchannel-doc
- qt5-qtwebchannel-examples
- qt5-qtwebsockets
- qt5-qtwebsockets-devel

- qt5-qtwebsockets-doc
- qt5-qtwebsockets-examples
- qt5-qtX11extras
- qt5-qtX11extras-devel
- qt5-qtX11extras-doc
- qt5-qtxmlpatterns
- qt5-qtxmlpatterns-devel
- qt5-qtxmlpatterns-doc
- qt5-qtxmlpatterns-examples
- qt5-rpm-macros
- qt5-srpm-macros
- webkit2gtk3
- webkit2gtk3-devel
- webkit2gtk3-jsc
- webkit2gtk3-jsc-devel
- wpebackend-fdo
- wpebackend-fdo-devel
- xorg-x11-server-Xorg

CHAPTER 11. KNOWN ISSUES

This part describes known issues in Red Hat Enterprise Linux 9.3.

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 installation program or use the **authselect** Kickstart command during installation.

Bugzilla:1640697^[1]

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^[1]

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. Instead, run Anaconda in a temporary virtual machine to keep the SELinux policy unchanged on a production system. 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 such as 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^[1]

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.

[Jira:RHEL-4707](#)

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.

[Jira:RHEL-4711](#)

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 **inobtcoun=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^[1]

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.

Jira:RHEL-4741

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, clear the **Connect automatically with priority** checkbox.
- e. Click **Save**.

[Bugzilla:2115783^{\[1\]}](#)

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 RM disk 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.

[Jira:RHEL-4762](#)

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](#)

Enabling the FIPS mode is not supported when building rpm-ostree images with RHEL image builder

Currently, there is no support to enable the FIPS mode when building **rpm-ostree** images with RHEL image builder.


[Jira:RHEL-4655](#)

Images built with the stig profile remediation fails to boot with FIPS error

FIPS mode is not supported by RHEL image builder. When using RHEL image builder customized with the **xccdf_org.ssgproject.content_profile_stig** profile remediation, the system fails to boot with the following error:

```
Warning: /boot//.vmlinuz-<kernel version>.x86_64.hmac does not exist
FATAL: FIPS integrity test failed
Refusing to continue
```

Enabling the FIPS policy manually after the system image installation with the **fips-mode-setup --enable** command does not work, because the **/boot** directory is on a different partition. System boots successfully if FIPS is disabled. Currently, there is no workaround available.



NOTE

You can manually enable FIPS after installing the image by using the **fips-mode-setup --enable** command.

[Jira:RHEL-4649](#)

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 stops responding, 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.

[Jira:RHEL-4737](#)

11.2. SECURITY

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](#)^[1]

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^{\[1\]}](#)

With a specific syntax, **scp** empties files copied to themselves

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.

[Jira:RHEL-1824](#)

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
```

- 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**.

[Jira:RHEL-1800](#)

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^{\[1\]}](#)

Keylime refuses runtime policies whose digests start with a backslash

The current script for generating runtime policies, **create_runtime_policy.sh**, uses SHA checksum functions, for example, **sha256sum**, to compute the file digest. However, when the input file name contains a backslash or **\n**, the checksum function adds a backslash before the digest in its output. In such cases, the generated policy file is malformed. When provided with the malformed policy file, the Keylime tenant produces the following or similar error message: **me.tenant - ERROR - Response code 400: Runtime policy is malformed**. To work around the problem, remove the backslash from the malformed policy file manually by entering the following command: **sed -i 's/^\\//g' <malformed_file_name>**.

[Jira:RHEL-11867^{\[1\]}](#)

Keylime agent rejects requests from the verifier after update

When the API version number of the Keylime agent (**keylime-agent-rust**) has been updated, the agent rejects requests that use a different version. As a consequence, if a Keylime agent is added to a verifier and then updated, the verifier tries to contact the agent using the old API version. The agent rejects this request and fails the attestation. To work around this problem, update the verifier (**keylime-verifier**) before updating the agent (**keylime-agent-rust**). As a result, when the agents are updated, the verifier detects the API change and updates its stored data accordingly.

[Jira:RHEL-1518^{\[1\]}](#)

The **fapolicyd** utility incorrectly allows executing changed files

Correctly, the IMA hash of a file should update after any change to the file, and **fapolicyd** should prevent execution of the changed file. However, this does not happen due to differences in IMA policy setup and in file hashing by the **evctml** utility. As a result, the IMA hash is not updated in the extended attribute of a changed file. Consequently, **fapolicyd** incorrectly allows the execution of the changed file.

Jira:RHEL-520^[1]

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 insecure SHA-1 signatures.

Bugzilla:2070722

OpenSCAP memory-consumption problems

On systems with limited memory, the OpenSCAP scanner might stop prematurely or it might not generate the results files. To work around this problem, you can customize the scanning profile to deselect rules that involve recursion over the entire / file system:

- **rpm_verify_hashes**
- **rpm_verify_permissions**
- **rpm_verify_ownership**
- **file_permissions_unauthorized_world_writable**
- **no_files_unowned_by_user**
- **dir_perms_world_writable_system_owned**
- **file_permissions_unauthorized_suid**
- **file_permissions_unauthorized_sgid**
- **file_permissions_ungroupowned**
- **dir_perms_world_writable_sticky_bits**

For more details and more workarounds, see the related [Knowledgebase article](#).

[Bugzilla:2161499](#)

Remediating service-related rules during kickstart installations might fail

During a kickstart installation, the OpenSCAP utility sometimes incorrectly shows that a service **enable** or **disable** state remediation is not needed. Consequently, OpenSCAP might set the services on the installed system to a non-compliant state. As a workaround, you can scan and remediate the system after the kickstart installation. This will fix the service-related issues.

[BZ#1834716](#)

11.3. RHEL FOR EDGE

The **open-vm-tools** package is not available in the **edge-vsphere** image

Currently, the **open-vm-tools** package is not installed by default in the **edge-vsphere** image. To workaround this issue, include the package in the blueprint customization. When using the **edge-vsphere** image type, add the **open-vm-tools** in the blueprint for the RHEL for Edge Container image or the RHEL for Edge Commit image.

[Jira:RHELDPCS-16574^{\[1\]}](#)

11.4. 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 work around this problem, restart the installation process.

[Bugzilla:2073510](#)

Running `createrepo_c` on local repositories generates duplicate `repodata` files

When you run the `createrepo_c` command on local repositories, it generates duplicate copies of `repodata` files, one of the copies is compressed and one is not. There is no workaround available, however, you can safely ignore the duplicate files. The `createrepo_c` command generates duplicate copies because of requirements and differences in other tools relying on repositories created by using `createrepo_c`.

[Bugzilla:2056318](#)

11.5. SHELLS AND COMMAND-LINE TOOLS

ReaR fails during recovery if the `TMPDIR` variable is set in the configuration file

Setting and exporting `TMPDIR` in the `/etc/rear/local.conf` or `/etc/rear/site.conf` ReaR configuration file does not work and is deprecated.

The ReaR default configuration file `/usr/share/rear/conf/default.conf` contains the following instructions:

```
# To have a specific working area directory prefix for Relax-and-Recover
# specify in /etc/rear/local.conf something like
#
# export TMPDIR="/prefix/for/rear/working/directory"
#
# where /prefix/for/rear/working/directory must already exist.
# This is useful for example when there is not sufficient free space
# in /tmp or $TMPDIR for the ISO image or even the backup archive.
```

The instructions mentioned above do not work correctly because the `TMPDIR` variable has the same value in the rescue environment, which is not correct if the directory specified in the `TMPDIR` variable does not exist in the rescue image.

As a consequence, setting and exporting `TMPDIR` in the `/etc/rear/local.conf` file leads to the following error when the rescue image is booted :

```
mktemp: failed to create file via template '/prefix/for/rear/working/directory/tmp.XXXXXXXXXX': No
such file or directory
cp: missing destination file operand after '/etc/rear/mappings/mac'
Try 'cp --help' for more information.
No network interface mapping is specified in /etc/rear/mappings/mac
```

or the following error and abort later, when running `rear recover`:

```
ERROR: Could not create build area
```

To work around this problem, if you want to have a custom temporary directory, specify a custom directory for ReaR temporary files by exporting the variable in the shell environment before executing ReaR. For example, execute the `export TMPDIR=...` statement and then execute the `rear` command in

the same shell session or script. As a result, the recovery is successful in the described configuration.

[Jira:RHEL-24847](#)

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](#)^[1]

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](#)^[1]

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
```

[Jira:RHEL-6105](#)

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?](#)

[Jira:RHEL-12009](#)

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.

Jira:RHEL-6995^[1]

11.6. 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, secondary 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](#)

11.7. NETWORKING

Using the XDP multi buffer mode with the **mlx5** driver and a MTU greater than 3498 bytes requires disabling RX Striding RQ

Running an eXpress Data Path (XDP) script with multi buffer mode on a host that matches all of the following conditions fails:

- The host uses the **mlx5** driver.
- The Maximum Transmission Unit (MTU) value is greater than 3498 bytes.
- The receive striding receive queue (RX Striding RQ) feature is enabled on the Mellanox interface.

If all conditions apply, the script fails with a **link set xdp fd failed** error. To run the XDP script on a host with a higher MTU, disable RX Striding RQ on the Mellanox interface:

```
# ethtool --set-priv-flags <interface_name> rx_striding_rq off
```

As a result, you can use the XDP multi buffer mode on interfaces that use the **mlx5** driver and have an MTU value greater than 3498 bytes.

Jira:RHEL-6496^[1]

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^[1]

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^[1]

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

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.

Jira:RHEL-9897^[1]

The Intel® i40e adapter permanently fails on IBM Power10

When the **i40e** adapter encounters an I/O error on IBM Power10 systems, the Enhanced I/O Error Handling (EEH) kernel services trigger the network driver's reset and recovery. However, EEH repeatedly reports I/O errors until the **i40e** driver reaches the predefined maximum of EEH freezes. As a consequence, EEH causes the device to fail permanently.

Jira:RHEL-15404^[1]

The **xdp-loader features** command fails

The **xdp-loader** utility was compiled against a previous version of **libbpf**. As a consequence, the **xdp-loader features** command fails with an error:

Cannot display features, because xdp-loader was compiled against an old version of libbpf without support for querying features.

No workaround is available. As a result, you cannot use the **xdp-loader features** command to display interface features.

Jira:RHEL-3382^[1]

11.8. 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^[1]

Customer applications with dependencies on kernel page size might 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 might 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 user space 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^{\[1\]}](#)

Upgrading to the latest real-time kernel with **dnf** does not install multiple kernel versions in parallel

Installing the latest real-time kernel with the **dnf** package manager requires resolving package dependencies to retain the new and current kernel versions simultaneously. By default, **dnf** removes the older **kernel-rt** package during the upgrade.

As a workaround, add the current **kernel-rt** package to the **installonlypkgs** option in the **/etc/yum.conf** configuration file, for example, **installonlypkgs=kernel-rt**.

The **installonlypkgs** option appends **kernel-rt** to the default list used by **dnf**. Packages listed in **installonlypkgs** directive are not removed automatically and therefore support multiple kernel versions to install simultaneously.

Note that having multiple kernels installed is a way to have a fallback option when working with a new kernel version.

[Bugzilla:2181571^{\[1\]}](#)

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 **iostat** 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^[1]

Hardware certification of the real-time kernel on systems with large core-counts might require passing the `skew-tick=1` boot parameter

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 displaying the kernel parameters you pass during boot.

```
cat /proc/cmdline
```

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.

Jira:RHEL-9318^[1]

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 on the LUKS-encrypted targets.

As a workaround, query the **Recommended crashkernel value** and gradually increase the memory size to an appropriate value. The **Recommended crashkernel value** can serve as reference to set the required memory size.

1. Print the estimate crash kernel 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.

Jira:RHEL-11196^[1]

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](#)

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 workaround is to power off the system and back on again. Do not reboot.

[Bugzilla:2129288^{\[1\]}](#)

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^{\[1\]}](#)

dkms provides an incorrect warning on program failure with correctly compiled drivers on 64-bit ARM CPUs

The Dynamic Kernel Module Support (**dkms**) utility does not recognize that the kernel headers for 64-bit ARM CPUs work for both the kernels with 4 kilobytes and 64 kilobytes page sizes. As a result, when the kernel update is performed and the **kernel-64k-devel** package is not installed, **dkms** provides an incorrect warning on why the program failed on correctly compiled drivers. To work around this problem, install the **kernel-headers** package, which contains header files for both types of ARM CPU architectures and is not specific to **dkms** and its requirements.

[JIRA:RHEL-25967^{\[1\]}](#)

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^[1]

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^[1]

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^[1]

Disabling quota accounting is no longer possible for an XFS filesystem mounted with quotas enabled

Starting with 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^[1]

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

NVMe/FC devices cannot be reliably used in a Kickstart file

NVMe/FC devices can be unavailable during parsing or execution of pre-scripts of the Kickstart file, which can cause the Kickstart installation to fail. To work around this issue, update the boot argument to **inst.wait_for_disks=30**. This option causes a delay of 30 seconds, and should provide enough time for

the NVMe/FC device to connect. With this workaround along with the NVMe/FC devices connecting in time, the Kickstart installation proceeds without issues.

[Jira:RHEL-8164^{\[1\]}](#)

Kernel panic while using the `qedi` driver

While using the `qedi` iSCSI driver, the kernel panics after OS boots. To work around this issue, disable the `kfence` runtime memory error detector feature by adding `kfence.sample_interval=0` to the kernel boot command line.

[Jira:RHEL-8466^{\[1\]}](#)

Unable to boot ARM based system with kernel-64k page size

While installing the `vdo` package, a kernel with 4k page size is installed as a dependency. As a consequence, the system boots with the 4k page size kernel even if you select 64k page size on the **Software Selection** screen. To work around this issue, select **Minimal Install** under **Base Environment** and 64k as page size under **Kernel options**. When the system boots for the first time, install additional softwares using the DNF package manager.

[Jira:RHEL-8354](#)

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. IDENTITY MANAGEMENT

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 \(PKINIT\)](#).

[Jira:RHEL-4902](#)

The DEFAULT:SHA1 subpolicy 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 subpolicy has to be set on RHEL 9 clients for PKINIT to work against AD KDCs](#).

[Jira:RHEL-4875](#)

FIPS support for AD trust requires the AD-SUPPORT crypto subpolicy

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** subpolicy 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.

[Jira:RHEL-4889](#)

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.

Jira:RHEL-12154^[1]

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.

Jira:RHEL-12143^[1]

Users without SIDs cannot log in to IdM after an upgrade

After upgrading your IdM replica to RHEL 9.2, the IdM Kerberos Distribution Center (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^[1]

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^[1]

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^{\[1\]}](#)

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.

[Jira:RHEL-4888](#)

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^[1]

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.

Jira:RHEL-4955

When the `nsslapd-numlisteners` attribute value is more than 2, Directory Server fails

If the `nsslapd-numlisteners` attribute value is higher than **2**, Directory Server might close the listening file descriptor instead of the accepted file descriptor. As a result, after some time, Directory Server stops listening on some ports and fails.

To work around the problem, set the `nsslapd-numlisteners` attribute value to **1**.

Jira:RHEL-17178^[1]

11.12. DESKTOP

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.

[Jira:RHEL-11924^{\[1\]}](#)

WebKitGTK fails to display web pages on IBM Z

The WebKitGTK web browser engine fails when trying to display web pages on the IBM Z architecture. The web page remains blank and the WebKitGTK process ends unexpectedly.

As a consequence, you cannot use certain features of applications that use WebKitGTK to display web pages, such as the following:

- The Evolution mail client
- The GNOME Online Accounts settings
- The GNOME Help application

[Jira:RHEL-4157](#)

11.13. 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^{\[1\]}](#)

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^{\[1\]}](#)

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^{\[1\]}](#)

11.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES

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.

[Jira:RHEL-1172](#)

11.15. 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 ends unexpectedly with an **internal error: process exited while connecting to monitor** message.

Similarly, using the RHEL 9 web console to install a guest operating system 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. 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^{\[1\]}](#)

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:207767^{\[1\]}](#)

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.

[Jira:RHEL-7337](#)

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.

[Jira:RHEL-7335](#)

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 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.

[Jira:RHEL-7336](#)

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^{\[1\]}](#)

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.

[Jira:RHEL-11366](#)

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 VMs to fail

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 hard-coded 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.

[Jira:RHEL-13335](#)^[1]

Redundant error messages on VMs with NVIDIA passthrough devices

When using an Intel host machine with a RHEL 9.2 and later 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](#)^[1]

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](#)^[1]

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.

[Jira:RHEL-333](#)

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.

[Jira:RHEL-7096](#)

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.

For more information, see the [Red Hat KnowledgeBase](#).

[Bugzilla:2073872](#)

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.

[Bugzilla:1970830](#)

Starting a VM with an NVIDIA A16 GPU sometimes causes the host GPU to stop working

Currently, if you start a VM that uses an NVIDIA A16 GPU passthrough device, the NVIDIA A16 GPU physical device on the host system in some cases stops working.

To work around the problem, reboot the hypervisor and set the **reset_method** for the GPU device to **bus**:

```
# echo bus > /sys/bus/pci/devices/<DEVICE-PCI-ADDRESS>/reset_method
# cat /sys/bus/pci/devices/<DEVICE-PCI-ADDRESS>/reset_method
bus
```

For details, see [the Red Hat Knowledgebase](#).

[Jira:RHEL-7212^{\[1\]}](#)

RT VMs with a FIFO scheduler cannot boot

Currently, after setting a real-time (RT) virtual machine (VM) to use the **fifo** setting for the vCPU scheduler, the VM becomes unresponsive when you attempt to boot it. Instead, the VM displays the **Guest has not initialized the display (yet)** error.

[Jira:RHEL-2815^{\[1\]}](#)

Windows VMs might become unresponsive due to storage errors

On virtual machines (VMs) that use Windows guest operating systems, the system in some cases becomes unresponsive when under high I/O load. When this happens, the system logs a **viostor Reset to device, \Device\RaidPort3, was issued** error.

[Jira:RHEL-1609^{\[1\]}](#)

Windows 10 VMs with certain PCI devices might become unresponsive on boot

Currently, a virtual machine (VM) that uses a Windows 10 guest operating system might become unresponsive during boot if a **virtio-win-scsi** PCI device with a local disk back end is attached to the VM. To work around the problem, boot the VM with the **multi_queue** option enabled.

[Jira:RHEL-1084^{\[1\]}](#)

The repair function of virtio-win-guest-tool for the virtio-win drivers does not work

Currently, when using the **Repair** button of **virtio-win-guest-tool** for a **virtio-win** driver, such as the Virtio Balloon Driver, the button has no effect. As a consequence, the driver cannot be reinstalled after being removed on the guest.

[Jira:RHEL-1517^{\[1\]}](#)

Windows 11 VMs with a memory balloon device set might close unexpectedly during reboot

Currently, rebooting virtual machines (VMs) that use a Windows 11 guest operating system and a memory balloon device in some cases fails with a **DRIVER POWER STAT FAILURE** blue-screen error.

[Jira:RHEL-935^{\[1\]}](#)

Migrating a Windows 11 or Windows Server 2022 VM under high network load sometimes fails

When live-migrating a virtual machine (VM) that uses Windows Server 2022 or Windows 11 as the guest operating system, the migration might become unresponsive or terminate unexpectedly if the network is impacted by high packet loss.

[Jira:RHEL-2316^{\[1\]}](#)

Resuming a postcopy VM migration fails in some cases

Currently, when performing a postcopy migration of a virtual machine (VM), if a proxy network failure occurs during the RECOVER phase of the migration, the VM becomes unresponsive and the migration cannot be resumed. Instead, the recovery command displays the following error:

```
error: Requested operation is not valid: QEMU reports migration is still running
```

[Jira:RHEL-7115](#)

The virtio balloon driver sometimes does not work on Windows 10 VMs

Under certain circumstances, the virtio-balloon driver does not work correctly on virtual machines (VMs) that use a Windows 10 guest operating system. As a consequence, such VMs might not use their assigned memory efficiently.

[Jira:RHEL-12118](#)

The virtio file system has suboptimal performance in Windows VMs

Currently, when a virtio file system (virtiofs) is configured on a virtual machine (VM) that uses a Windows guest operating system, the performance of virtiofs in the VM is significantly worse than in VMs that use Linux guests.

[Jira:RHEL-1212^{\[1\]}](#)

Hot-unplugging a storage device on Windows VMs might fail

On virtual machines (VMs) that use a Windows guest operating system, removing a storage device when the VM is running (also known as a device hot-unplug) in some cases fails. As a consequence, the storage device remains attached to the VM and the disk manager service might become unresponsive.

[Jira:RHEL-869](#)

Hot plugging CPUs to a Windows VM might cause a system failure

When hot plugging the maximum number of CPUs to a Windows virtual machine (VM) with huge pages enabled, the guest operating system might crash with the following *Stop error*:

```
PROCESSOR_START_TIMEOUT
```

[Jira:RHEL-1220](#)

Updating virtio drivers on Windows VMs might fail

When updating the KVM paravirtualized (**virtio**) drivers on a Windows virtual machine (VM), the update might cause the mouse to stop working and the newly installed drivers might not be signed. This problem occurs when updating the **virtio** drivers by installing from the **virtio-win-guest-tools** package, which is a part of the **virtio-win.iso** file.

To work around this problem, update the **virtio** drivers by using Windows Device Manager.

[Jira:RHEL-574^{\[1\]}](#)

Kdump fails on virtual machines with AMD SEV-SNP

Currently, kdump fails on RHEL 9 virtual machines (VMs) that use the AMD Secure Encrypted Virtualization (SEV) with the Secure Nested Paging (SNP) feature.

[Jira:RHEL-10019^{\[1\]}](#)

11.16. 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. Re-create 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^[1]

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^[1]

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^[1]

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 Knowledge Base](#).

Jira:RHEL-12122

Large VMs might fail to boot into the debug kernel when the **kmemleak** option is enabled

When attempting to boot a RHEL 9 virtual machine (VM) into the debug kernel, the booting might fail with the following error if the machine kernel is using the **kmemleak=on** argument.

```
Cannot open access to console, the root account is locked.  
See sulin(8) man page for more details.
```

```
Press Enter to continue.
```

This problem affects mainly large VMs because they spend more time in the boot sequence.

To work around the problem, edit the `/etc/fstab` file on the machine and add extra timeout options to the `/boot` and `/boot/efi` mount points. For example:

```
UUID=e43ead51-b364-419e-92fc-b1f363f19e49 /boot xfs defaults,x-systemd.device-timeout=600,x-systemd.mount-timeout=600 0 0

UUID=7B77-95E7 /boot/efi vfat defaults,uid=0,gid=0,umask=077,shortname=winnt,x-systemd.device-timeout=600,x-systemd.mount-timeout=600 0 2
```

Jira:RHELDOCS-16979^[1]

11.17. 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^[1]

11.18. 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^[1]

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:2188627 , Bugzilla:1987471 , Bugzilla:2149025 , Bugzilla:2166332 , Bugzilla:2189946 , Bugzilla:2189954 , Bugzilla:1975930 , Bugzilla:1974242 , Bugzilla:1759941 , Bugzilla:2053204 , Bugzilla:2116948 , Bugzilla:2179278 , Bugzilla:2189717 , Bugzilla:2170494 , Bugzilla:2098236 , Jira:RHEL-17178
NetworkManager	Bugzilla:2176137 , Bugzilla:2161915 , Bugzilla:2151986 , Bugzilla:2190375 , Bugzilla:2069001 , Bugzilla:2069004 , Bugzilla:2148684 , Bugzilla:2158328 , Bugzilla:2180966 , Bugzilla:2151040 , Bugzilla:1894877
Release Notes	Jira:RHELDOCS-16861 , Jira:RHELDOCS-16760 , Jira:RHELDOCS-16756 , Jira:RHELDOCS-16612 , Jira:RHELDOCS-17102 , Jira:RHELDOCS-16979
anaconda	Bugzilla:2171811 , Bugzilla:2164819 , Bugzilla:2177219 , Bugzilla:2157921 , Bugzilla:2065754 , Bugzilla:2107346 , Bugzilla:2127473 , Bugzilla:2050140 , Bugzilla:1877697 , Jira:RHEL-4707 , Jira:RHEL-4711 , Bugzilla:1997832 , Jira:RHEL-4741 , Bugzilla:2115783 , Jira:RHEL-4762 , Bugzilla:2163497 , Jira:RHEL-4737
ansible-freeipa	Bugzilla:2175767 , Bugzilla:2127903 , Bugzilla:2127907
audit	Jira:RHELPLAN-161087
bacula	Jira:RHEL-6856
bind	Bugzilla:1984982
cloud-init	Bugzilla:2118235 , Bugzilla:2172341 , Jira:RHEL-12122
cockpit	Bugzilla:2203361
cockpit-appstream	Bugzilla:2030836
cockpit-machines	Bugzilla:2173584
crash	Bugzilla:2170283
createrepo_c	Bugzilla:2056318
crypto-policies	Bugzilla:2216257 , Bugzilla:2193324 , Jira:RHEL-591 , Bugzilla:2225222
cups-filters	Bugzilla:2229784

Component	Tickets
cyrus-sasl	Bugzilla:1995600
debugedit	Bugzilla:2177302
device-mapper-multipath	Jira:RHEL-782 , Bugzilla:2164869 , Bugzilla:2033080 , Bugzilla:2011699 , Bugzilla:1926147
device-mapper-persistent-data	Bugzilla:2175198
dnf	Bugzilla:2124793 , Bugzilla:2212262 , Bugzilla:2073510
dnf-plugins-core	Bugzilla:2157844 , Bugzilla:2134638 , Bugzilla:2203100
edk2	Bugzilla:1935497
elfutils	Bugzilla:2182061 , Bugzilla:2182059
fapolicyd	Jira:RHEL-624 , Jira:RHEL-622 , Jira:RHEL-817 , Bugzilla:2054740 , Jira:RHEL-520
fence-agents	Bugzilla:2187327
fuse3	Bugzilla:2188182
gcc	Bugzilla:2193180 , Bugzilla:2168204 , Bugzilla:2208908
gcc-toolset-13	Bugzilla:2171919
gcc-toolset-13-annobin	Bugzilla:2171923
gcc-toolset-13-binutils	Bugzilla:2171926
gcc-toolset-13-gcc	Bugzilla:2172093
gcc-toolset-13-gdb	Bugzilla:2172096
gfs2-utils	Bugzilla:2170017
gimp	Bugzilla:2047161
glibc	Bugzilla:2169978 , Bugzilla:2213907 , Bugzilla:2177235
gnupg2	Bugzilla:2073567 , Bugzilla:2070722

Component	Tickets
gnutls	Bugzilla:2157953 , Bugzilla:2108532
golang	Bugzilla:2185259 , Bugzilla:2111072 , Bugzilla:2092016
grafana	Bugzilla:2193018 , Bugzilla:2190025
grub2	Bugzilla:2184069
gssproxy	Bugzilla:2181465
gtk3	Jira:RHEL-11924
httpd	Bugzilla:2184403 , Bugzilla:2173295
ipa	Bugzilla:2196426 , Bugzilla:2165880 , Bugzilla:2229712 , Bugzilla:2227831 , Bugzilla:2084180 , Bugzilla:2084166 , Bugzilla:2069202 , Bugzilla:2094673 , Bugzilla:2057471 , Jira:RHEL-12154 , Jira:RHEL-12143 , Jira:RHEL-4955
iproute	Jira:RHEL-428
java-17-openjdk	Bugzilla:2186647
jmc-core	Bugzilla:1980981
kdump-anaconda-addon	Jira:RHEL-11196
kernel	Bugzilla:1898184 , Bugzilla:2177180 , Bugzilla:2144528 , Bugzilla:2210263 , Bugzilla:2180124 , Bugzilla:2192730 , Bugzilla:2178741 , Bugzilla:2195986 , Bugzilla:2208365 , Bugzilla:2187856 , Bugzilla:2192722 , Bugzilla:2171093 , Bugzilla:2189292 , Bugzilla:2193330 , Bugzilla:2178930 , Bugzilla:2092194 , Bugzilla:2101598 , Bugzilla:2218207 , Bugzilla:2173947 , Bugzilla:2178956 , Bugzilla:2173594 , Bugzilla:1613522 , Bugzilla:1874182 , Bugzilla:1995338 , Bugzilla:1570255 , Bugzilla:2177256 , Bugzilla:2178699 , Bugzilla:2023416 , Bugzilla:2021672 , Bugzilla:2027304 , Bugzilla:1660337 , Bugzilla:1955275 , Bugzilla:2142102 , Bugzilla:2068237 , Bugzilla:2040643 , Bugzilla:2186375 , Bugzilla:2183538 , Bugzilla:2206599 , Bugzilla:2167783 , Bugzilla:2000616 , Bugzilla:2013650 , Bugzilla:2132480 , Bugzilla:2059545 , Bugzilla:2005173 , Bugzilla:2128610 , Bugzilla:2129288 , Bugzilla:2013884 , Bugzilla:2149989
kernel / Networking / IPSec	Jira:RHEL-1015
kernel / Networking / NIC Drivers	Jira:RHEL-6496 , Jira:RHEL-9897 , Jira:RHEL-15404
kernel / Platform Enablement / NVMe	Jira:RHEL-8171 , Jira:RHEL-8164

Component	Tickets
kernel / Storage / Storage Drivers	Jira:RHEL-8466
kernel / Virtualization / KVM	Jira:RHEL-7212 , Jira:RHEL-2815
kernel-rt	Bugzilla:2181571
kernel-rt / Other	Jira:RHEL-9318
kexec-tools	Bugzilla:2083475 , Bugzilla:2173815 , Bugzilla:2169720 , Bugzilla:2160676 , Bugzilla:2113873 , Bugzilla:2064708
keylime	Jira:RHEL-595 , Jira:RHEL-11866 , Jira:RHEL-392 , Jira:RHEL-393 , Jira:RHEL-947 , Jira:RHEL-1252 , Jira:RHEL-11867 , Jira:RHEL-1518
keylime-agent-rust	Jira:RHEL-476 , Jira:RHEL-395 , Jira:RHEL-396
kmod	Bugzilla:2103605
kmod-kvdo	Jira:RHEL-8354
krb5	Bugzilla:2178298 , Bugzilla:2155607 , Jira:RHEL-4902 , Bugzilla:2060798 , Jira:RHEL-4875 , Jira:RHEL-4889 , Bugzilla:2060421 , Bugzilla:2016312 , Jira:RHEL-4888
libabigail	Bugzilla:2186931
libotr	Bugzilla:2086562
libpfm	Bugzilla:2185652
libvirt	Bugzilla:2032406 , Bugzilla:2168499 , Bugzilla:2014487 , Bugzilla:2143158 , Bugzilla:2078693
libxcrypt	Bugzilla:2034569
llvm-toolset	Bugzilla:2178796
lvm2	Bugzilla:2038183
mysql	Bugzilla:1991500
nfs-utils	Bugzilla:2081114
nginx-1.22-module	Bugzilla:2170808

Component	Tickets
nmstate	Bugzilla:2179916 , Bugzilla:2180795 , Bugzilla:2177733 , Bugzilla:2183214 , Bugzilla:2187622
nodejs	Bugzilla:2186717
nss	Bugzilla:2157950
nvme-cli	Bugzilla:2159929
nvme-stas	Bugzilla:1893841
open-vm-tools	Bugzilla:2037657
opencryptoki	Bugzilla:2160061
opensc	Jira:RHEL-280
openscap	Bugzilla:2217442 , Bugzilla:2161499
openslp	Jira:RHEL-6995
openssh	Bugzilla:2070163 , Bugzilla:2056884
openssl	Bugzilla:2216256 , Bugzilla:2153471 , Bugzilla:2188180 , Bugzilla:2160797 , Bugzilla:2168665 , Bugzilla:1975836 , Bugzilla:1681178 , Bugzilla:1685470
osbuild	Jira:RHEL-4655
osbuild-composer	Bugzilla:2173928 , Jira:RHEL-7999 , Jira:RHEL-4649
oscap-anaconda-addon	Bugzilla:2172264 , Jira:RHEL-1824
pacemaker	Bugzilla:2189301 , Bugzilla:2182482
papi	Bugzilla:2111923 , Bugzilla:2186927 , Bugzilla:2215582
pause-container	Bugzilla:2106816
pcp	Bugzilla:2175602 , Bugzilla:2185803
pcs	Bugzilla:2168155 , Bugzilla:2163953 , Bugzilla:2175881 , Bugzilla:2182810 , Bugzilla:1423473 , Bugzilla:2177996 , Bugzilla:1860626 , Bugzilla:2163914
pcsc-lite-ccid	Bugzilla:2209457

Component	Tickets
perl-HTTP-Tiny	Bugzilla:2228412
pki-core	Bugzilla:2084181
podman	Jira:RHELPLAN-154314 , Jira:RHELPLAN-154432 , Jira:RHELPLAN-154441 , Jira:RHELPLAN-154438 , Jira:RHELPLAN-163003 , Jira:RHELPLAN-160660 , Jira:RHELPLAN-154429 , Bugzilla:2069279
postfix	Bugzilla:2134789
python-greenlet	Bugzilla:2149497
python3.11-lxml	Bugzilla:2157708
qemu-kvm	Bugzilla:1880531 , Bugzilla:1965079 , Bugzilla:1951814 , Bugzilla:2060839 , Bugzilla:2014229 , Jira:RHEL-7335 , Jira:RHEL-7336 , Bugzilla:1915715 , Jira:RHEL-13335 , Jira:RHEL-333 , Bugzilla:2176010 , Bugzilla:2058982 , Bugzilla:2073872
qemu-kvm / Devices	Jira:RHEL-1220
qemu-kvm / Graphics	Jira:RHEL-7135
qemu-kvm / Live Migration	Jira:RHEL-7096 , Jira:RHEL-2316 , Jira:RHEL-7115
qemu-kvm / Networking	Jira:RHEL-7337
rear	Bugzilla:2188593 , Bugzilla:2172912 , Bugzilla:2196445 , Bugzilla:2145014
redis	Bugzilla:2129826
resource-agents	Bugzilla:2174911 , Bugzilla:2142518 , Bugzilla:2142002 , Bugzilla:2182415 , Bugzilla:2179003
restore	Bugzilla:1997366
rhel-system-roles	Bugzilla:2224384 , Bugzilla:2216753 , Bugzilla:2224385 , Bugzilla:2185065 , Bugzilla:2181656 , Bugzilla:2211194 , Bugzilla:2218592 , Bugzilla:2211723 , Bugzilla:2218204 , Bugzilla:2151373 , Bugzilla:2179460 , Bugzilla:2211748 , Bugzilla:2229802 , Bugzilla:2181657 , Bugzilla:2168692 , Bugzilla:2211984 , Bugzilla:2232241 , Bugzilla:2232231 , Bugzilla:2224090 , Bugzilla:2222761 , Bugzilla:2223764 , Bugzilla:2222428 , Bugzilla:2216520 , Bugzilla:2211187 , Bugzilla:2209200 , Bugzilla:2193058 , Bugzilla:2186057 , Jira:RHEL-1499 , Jira:RHEL-1397 , Jira:RHEL-906 , Jira:RHEL-1495 , Jira:RHEL-898 , Jira:RHEL-885 , Bugzilla:1999770 , Bugzilla:2123859 , Jira:RHEL-1172 , Bugzilla:2186218

Component	Tickets
rpm	Bugzilla:2157836
rsyslog	Jira:RHELPLAN-160541
rust	Bugzilla:2191743 , Bugzilla:2227082
s390utils	Bugzilla:1932480
samba	Bugzilla:2190415
scap-security-guide	Bugzilla:2221697 , Bugzilla:2155790 , Jira:RHEL-1905 , Bugzilla:2203791 , Bugzilla:2213958 , Bugzilla:2223178 , Bugzilla:2193169 , Jira:RHEL-1800 , Bugzilla:2038978
selinux-policy	Bugzilla:2080443 , Bugzilla:2170495 , Bugzilla:2184999 , Bugzilla:2162663 , Bugzilla:2112729 , Jira:RHELPLAN-163014 , Bugzilla:2187745 , Bugzilla:2229722 , Bugzilla:2064274
setools	Bugzilla:2231801
sevctl	Bugzilla:2104857
sos	Bugzilla:1869561
squid-container	Bugzilla:2178953
sssd	Bugzilla:2065693 , Bugzilla:2056482 , Bugzilla:1608496
stratisd	Bugzilla:2041558
subscription-manager	Bugzilla:2163716 , Bugzilla:2136694
sysstat	Jira:RHEL-12009
systemd	Bugzilla:2018112 , Jira:RHEL-6105
systemtap	Bugzilla:2186934
tang	Bugzilla:2188743
tigervnc	Bugzilla:2060308
tuned	Bugzilla:2113900

Component	Tickets
ubi9-micro-container	Bugzilla:2223028
udisks2	Bugzilla:1983602 , Bugzilla:2213769
unbound	Bugzilla:2070495
valgrind	Bugzilla:2124346
virt-v2v	Bugzilla:2168082
virtio-win	Bugzilla:1969724 , Jira:RHEL-11366 , Jira:RHEL-910 , Jira:RHEL-1609 , Jira:RHEL-869
virtio-win / distribution	Jira:RHEL-1517 , Jira:RHEL-574
virtio-win / virtio-win-prewhql	Jira:RHEL-1084 , Jira:RHEL-935 , Jira:RHEL-12118 , Jira:RHEL-1212
webkit2gtk3	Jira:RHEL-4157
which	Bugzilla:2181974
xdp-tools	Bugzilla:2218500 , Jira:RHEL-3382

Component	Tickets
other	<p>Bugzilla:2232554, Jira:RHELDOCS-17055, Jira:RHELPLAN-163133, Jira:RHELPLAN-163665, Jira:RHELDOCS-16405, Jira:RHELDOCS-16247, Bugzilla:2136937, Jira:RHELDOCS-16474, Jira:RHELDOCS-16462, Jira:RHELDOCS-16386, Jira:RHELPLAN-156196, Jira:RHELDOCS-16708, Jira:RHELDOCS-16709, Jira:RHELDOCS-16339, Jira:RHELDOCS-16877, Jira:RHELPLAN-122345, Jira:RHELDOCS-16487, Jira:RHELDOCS-16752, Jira:RHELDOCS-17101, Bugzilla:2236182, Jira:RHELDOCS-17040, Bugzilla:2020529, Bugzilla:2030412, Jira:RHELPLAN-103993, Jira:RHELPLAN-27394, Jira:RHELPLAN-27737, Jira:RHELDOCS-16861, Jira:RHELDOCS-17050, Bugzilla:1927780, Jira:RHELPLAN-110763, Bugzilla:1935544, Bugzilla:2089200, Jira:RHELDOCS-16948, Jira:RHELPLAN-99136, Jira:RHELDOCS-17380, 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-139805, Jira:RHELDOCS-16756, Jira:RHELPLAN-153267, Jira:RHELDOCS-16300, Jira:RHELDOCS-16432, Jira:RHELDOCS-16393, Jira:RHELDOCS-16612, Jira:RHELDOCS-17102, 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, Bugzilla:1970830, Jira:RHELDOCS-16574</p>

APPENDIX B. REVISION HISTORY

0.2-2

Tue Jun 11 2024, Brian Angelica (bangelic@redhat.com)

- Add Deprecated Functionality [RHELDOCS-18049](#) (Shells and command-line tools).

0.2-1

Tue Jun 11 2024, Brian Angelica (bangelic@redhat.com)

- Added an Known Issue [RHEL-24847](#) (Shells and command-line tools).

0.2-0

Thu May 16 2024, Gabriela Fialová (gfialova@redhat.com)

- Added an Known Issue [RHEL-10019](#) (Virtualization).

0.1-9

Thu Apr 18 2024, Gabriela Fialová (gfialova@redhat.com)

- Added an Enhancement [RHEL-19142](#) (Networking).

0.1-8

Thu Apr 11 2024, Gabriela Fialová (gfialova@redhat.com)

- Added an Enhancement [BZ#1513934](#) (IdM)

0.1-7

Thu Mar 14 2024, Gabriela Fialová (gfialova@redhat.com)

- Added a Known Issue [JIRA:RHEL-25967](#) (Kernel)

0.1-6

Mon Mar 04 2024, Gabriela Fialová (gfialova@redhat.com)

- Added a bug fix [Jira:SSSD-6096](#) (Identity Management).

0.1-5

Wed Feb 28 2024, Gabriela Fialová (gfialova@redhat.com)

- Updated a Known Issue to Bug Fix [RHEL-8171](#) (Storage).

0.1-4

Wed Feb 7 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added a new feature [RHEL-14694](#) (Networking).

0.1-3

Thu Feb 1 2024, Gabriela Fialová (gfialova@redhat.com)

- Added a KI [BZ#1834716](#) (Security)

- Updated s Deprecated Functionality [RHELDOCS-16756](#) (Container tools)

0.1-2

Mon Jan 29 2024, Gabriela Fialová (gfialova@redhat.com)

- Added a bug fix [RHELPLAN-157337](#) (Security)

0.1-1

Thu Jan 2024, Lenka Špačková (lspackova@redhat.com)

- Added an enhancement related to Python [RHELDOCS-17369](#) (Dynamic programming languages, web and database servers)

0.1-0

Wed Jan 10 2024, Gabriela Fialová (gfialova@redhat.com)

- Add Deprecated Functionality [RHELDOCS-17380](#) (Security)

0.0-9

Tue Jan 2 2024, Gabriela Fialová (gfialova@redhat.com)

- Updated description in Enhancement [BZ#2184403](#)

0.0-8

Thu Nov 23 2023, Gabriela Fialová (gfialova@redhat.com)

- Added KI [RHEL-8354](#) (installer)

0.0-7

Wed Nov 22 2023, Gabriela Fialová (gfialova@redhat.com)

- Add IdM KI [RHEL-17178](#)

0.0-6

Tue Nov 21 2023, David Vozenilek (dvozenil@redhat.com)

- Add System Roles RNs [BZ#2211723](#), [BZ#2218204](#), [BZ#2186057](#)

0.0-5

Mon Nov 20 2023, Jana Heves (jsvarova@redhat.com)

- Add KI [RHEL-15404](#) sst_kernel_generalists

0.0-4

Sun Nov 19 2023, Filip Hanzelka (fhanzelk@redhat.com)

- Add BF in IdM [RHELDOCS-17011](#)

0.0-3

Thu Nov 16 2023, Marek Suchánek (msuchane@redhat.com)

- Deprecate Inkscape and LibreOffice Flatpak [RHELDOCS-17102](#)

0.0-2

Thu Nov 16 2023, Lenka Špačková (lspackova@redhat.com)

- **Node.js 20** is now fully supported ([BZ#2186717](#)).

0.0-1

Wed Nov 08 2023, Gabriela Fialová (gfialova@redhat.com)

- Release of the Red Hat Enterprise Linux 9.3 Release Notes.

0.0-0

Wed Sep 27 2023, Gabriela Fialová (gfialova@redhat.com)

- Release of the Red Hat Enterprise Linux 9.3 Beta Release Notes.