Red Hat Enterprise Linux 9

9.7 Release Notes

Release Notes for Red Hat Enterprise Linux 9.7

Last Updated: 2025-11-12

Red Hat Enterprise Linux 9 9.7 Release Notes

Release Notes for Red Hat Enterprise Linux 9.7

Legal Notice

Copyright © 2025 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 9.7 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

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 9.7

Installer and image creation

Key highlights for RHEL image builder:

- You can use RHEL image builder to create disk images with advanced partitioning.
- You can customize your blueprint to enable injecting a Kickstart file when building ISO images.
- System images created with the RHEL image builder, such as AWS or KVM formats, do not have a separate /**boot** partition.
- RHEL Image Builder now supports WSL2 images.

For more information, see New features - Installer and image creation .

Security

The system-wide **cryptographic policies** introduce the PQ subpolicy to enable post-quantum cryptography (PQC) algorithms and contain many improvements to support PQC in applications.

OpenSSL 3.5 introduces support for the ML-KEM, ML-DSA, and SLH-DSA **post-quantum algorithms** and adds the hybrid ML-KEM algorithms to the default TLS group list.

See New features - Security for more information.

Kernel

RHEL 9.7 adds kernel-area improvements that advance performance analysis and energy telemetry. Updates include a refined jitter entropy source in **rng-tools**, upstream alignment for **perf** and BPF to v6.14 and v6.15, expanded **uncore** and **core** event counters, Intel RAPL energy events, Intel Trace Hub (NPK) device IDs, and refreshed crash analysis and **python-drgn** tooling. NVMe-TCP **kdump** to an NVMe namespace can fail in some environments. Apply operational caution where this limitation is relevant.

Dynamic programming languages, web and database servers

Later versions of the following Application Streams are now available:

- **Node.js 24**

See New features - Dynamic programming languages, web and database servers and Technology Previews - Dynamic programming languages, web and database servers for more information.

Compilers and development tools

Updated system toolchain

The following system toolchain components have been updated:

- **Glibc 2.34**
- **Annobin 12.98**

Updated performance tools and debuggers

The following performance tools and debuggers have been updated in RHEL 9.7:

- **GDB 16.3**

- **Valgrind 3.25.1**
- **SystemTap 5.3**
- **Dyninst 13.0.0**
- **elfutils 0.193**
- **libabigail 2.8**

Updated performance monitoring tools
The following performance monitoring tools have been updated in RHEL 9.7:

- **PCP 6.3.7**
- **Grafana 10.2.6**

.NET 10.0 is now available on RHEL
Red Hat Enterprise Linux (RHEL) supports .NET, a general-purpose development platform that features automatic memory management and modern programming languages, allowing you to build high-quality applications efficiently. This update adds support for the most recent version, .NET 10.0 (Long-Term Support), expanding the versions available on RHEL. Other supported versions include .NET 9.0 (Standard-Term Support) and the previous long-term support version, .NET 8.0.

For more information, see Release Notes for .NET 10.0 RPM packages and Release Notes for .NET 10.0 containers

Updated compiler toolsets
The following compiler toolsets have been updated in RHEL 9.7:

- **GCC Toolset 15**
  - **GCC 15.1**
  - **Binutils 2.44**
    Note that **Annobin** and **dwz** are not provided in GCC Toolset starting with version 15.
- **LLVM Toolset 20.1.8**
- **Rust Toolset 1.88.0**
- **Go Toolset 1.24**

For detailed changes, see New features - Compilers and development tools .

The web console
The **cockpit** packages have been upgraded to version 344, which provides many improvements, most notably the upgrade to the Patternfly 6 system design.

See New features - The web console for more information.

1.2. IN-PLACE UPGRADE

In-place upgrade from RHEL 8 to RHEL 9
The supported in-place upgrade paths currently are:

- From RHEL 8.10 to RHEL 9.4, 9.6, and 9.7 on the following architectures:

    - 64-bit Intel, AMD, and ARM
    - IBM POWER 9 (little endian) and later
    - IBM Z architectures, excluding z13
- From RHEL 8.10 to RHEL 9.4 and 9.6 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 .

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 and bug fixes include:

- Fix in-place upgrades on systems that use the **fapolicyd** software framework.
- Disable the **localpkg_gpgcheck** DNF option when performing the upgrade allowing the required installation of bundled **leapp-deps-el10** and **leapp-repository-deps-el10** metapackages.
- Introduce the LiveMode feature as a Technology Preview. LiveMode allows you to upgrade by using the standard booting process. You can also use LiveMode for troubleshooting and testing. For more information, see Configuring the upgrade with LiveMode.
- Inhibit the upgrade on systems that use deprecated **network-legacy** dracut module to prevent kernel panic.
- Migrate SSSD configuration during the in-place upgrade.
- Enable upgrades on PAYG RHEL systems that use Red Hat Upgrade Infrastructure (RHUI) on Amazon Web Services (AWS), Azure, and Google Cloud.

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 In-place upgrades over multiple RHEL major versions by using Leapp.

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

- VNC Configurator
- Red Hat Satellite Upgrade Helper
- JVM Options Configuration Tool
- Load Balancer Configuration Tool
- Ceph Placement Groups (PGs) per Pool Calculator
- Yum Repository Configuration Helper
- Red Hat Out of Memory Analyzer
- Postfix Configuration Helper
- System Unit Generator
- Rsyslog Configuration Helper
- Red Hat IdM Upgrade Helper

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 .

Red Hat Insights is now **Red Hat Lightspeed**. This is a change in name only and all the same product features, functionalities, and capabilities you have relied on under the Red Hat Insights name remain under the name Red Hat Lightspeed. With Red Hat Lightspeed, which is included with all RHEL subscriptions, you can proactively identify, examine, and resolve known technical issues. For instructions on how to install the client and register your system to the service, see the Red Hat Lightspeed documentation page.

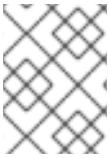

NOTE

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 9.7 is distributed with the kernel version 5.14.0-611.5.1, 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 Interactively installing RHEL from installation media 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 Automatically installing RHEL 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.

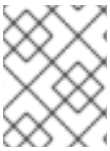

NOTE

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 9.7.

4.1. INSTALLER AND IMAGE CREATION

New boot menu entry for fips=1 added to ISO installations

With this update, the DVD and Boot ISO image installations provide a new boot menu entry for setting the **fips=1** kernel boot option. This simplifies the process, as enabling FIPS mode during the RHEL installation ensures that the system generates all keys with FIPS-approved algorithms and continuous monitoring tests in place. By using this boot option, you start the installation with the **fips=1** kernel parameter and you can target the system's compliance with Federal Information Processing Standards (FIPS) 140 requirements.

Jira:RHEL-91930

The blueprint file customization now supports a URI field for referencing files from external sources

This update adds the **URI** field support to the blueprint file customization structure. As a result, you can reference and source files from external locations rather than only those included directly in the blueprint, providing more flexible customization of the build system and a more adaptable build experience.

Jira:RHELDOCS-21016[1]

RHEL image builder supports a new image type vagrant-libvirt for vagrant

With this update, RHEL image builder supports the **libvirt** hypervisor, and you can easily run RHEL virtual machines by using Vagrant. This enhancement provides pre-configured images to ensure a consistent and streamlined setup. It also grants sudo privileges to the **vagrant** user within the Vagrant box, making it easier to manage and execute administrative tasks. These enhancements deliver a more efficient and seamless experience when working with RHEL virtual machines in Vagrant environments.

Jira:RHELDOCS-21025[1]

RHEL Image Builder GUI supports modularized content discovery

Starting from RHEL 9.7, RHEL Image Builder Graphical User Interface (GUI) supports modularized content discovery. This capability introduces the following enhancements:

- When creating RHEL OS images, you can use the RHEL Image Builder GUI to discover and include modularized content from various repositories, including RHEL AppStream and third-party repositories, for example, Extra Packages for Enterprise Linux (EPEL).
- Enhanced modularity support in RHEL. Application Streams leverage DNF modularity and **modulemd** metadata to provide flexible package management. You can specify version streams and use case profiles in the modules with support for default streams and profiles.
- DNF modularity implementation updates. The @ character syntax for specifying RPM groups enables and installs module streams, providing compatibility for kickstart files.

Jira:RHELDOCS-21026[1]

RHEL Image Builder now supports WSL2 images

You can now use the RHEL image builder to create Windows Subsystem for Linux (WSL2). The image type is available in the **wsl** format, and to consume the image, deploy it by double-clicking the generated file.

Jira:RHELDOCS-20633[1]

A new rhel9/bootc-image-builder container image is generally available in RHEL

The rhel9/bootc-image-builder container image for image mode for RHEL includes a minimal version of image builder that converts bootable container images, for example rhel-bootc, to different disk image formats, such as QCOW2, AMI, VMDK, ISO, and others.

Jira:RHELDOCS-17733[1]

The bootc-image-builder tool is generally available in RHEL

The **bootc-image-builder** tool, now generally available in RHEL, works as a container to easily create and deploy compatible disk images from the **bootc** container inputs. After running your container image with **bootc-image-builder**, you can generate images for the architecture that you need. Then, you can deploy the resulting image on VMs, clouds, or servers. You can easily update the images with the bootc, instead of having to regenerate the content with **bootc-image-builder** every time a new update is required.

Jira:RHELDOCS-17468[1]

composefs read-only file system supports bootc/ostree and podman projects

The composefs read-only file system is generally intended only to be used by the bootc/ostree and podman projects at the current time. With composefs, you can use these projects to create and use read-only images, share file data between images, and validate images at runtime. As a result, you have a fully verified file-system tree mounted, with opportunistic fine-grained sharing of identical files.

Jira:RHEL-18157[1]

4.2. SECURITY

NSS rebased to 3.112

The NSS cryptographic toolkit packages have been rebased to upstream version 3.112, which provides many improvements and fixes. Most notably, the following:

- Added support for the Module-Lattice-Based Digital Signature Algorithm (ML-DSA), which is a post-quantum cryptography (PQC) standard.
- Added hybrid support for SSL for the MLKEM1024 key encapsulation mechanism.

The following known issues occur in this version:

- Updating the NSS database password corrupts the ML-DSA seed. For more information, see RHEL-114443.

Jira:RHEL-103366

RHEL 9.7 crypto-policies supports post-quantum cryptography

With this update of the system-wide cryptographic policies, you can enable support for post-quantum cryptography (PQC) through the new PQ subpolicy. The most notable changes in RHEL 9.7 **crypto-policies** include:

- After you apply the PQ subpolicy, for example, by using the **update-crypto-policies --set DEFAULT:PQ** command, hybrid Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) and pure Module-Lattice-Based Digital Signature Standard (ML-DSA) post-quantum cryptographic algorithms are enabled in LEGACY, DEFAULT, FUTURE, and FIPS cryptographic policies with the highest priorities.
- The PQC algorithms are enabled for the Sequoia PGP tool in all policies with the PQ subpolicy.
- The new OpenSSL group selection syntax prioritizes post-quantum groups over classical ones if you enable the PQ subpolicy. You can revert this behavior only by disabling all PQ groups.
- The ML-DSA-44, ML-DSA-65, and ML-DSA-87 PQC algorithms are enabled for NSS TLS connections in all cryptographic policies with the PQ subpolicy.
- The PQ subpolicy also enables the **mlkem768x25519**, **secp256r1mlkem768**, and **secp384r1mlkem1024** hybrid ML-KEM groups for NSS TLS negotiations.

Jira:RHEL-91839, Jira:RHEL-103963, Jira:RHEL-106866, Jira:RHEL-103786, Jira:RHEL-97764

OpenSSL rebased to 3.5

OpenSSL is rebased to upstream version 3.5. This version provides important fixes and enhancements, most notably the following:

- Added support for the ML-KEM, ML-DSA, and SLH-DSA post-quantum algorithms.
- Added the hybrid ML-KEM algorithms to the default TLS group list.
- Enhanced TLS configuration options.
- Added support for the QUIC transport protocol according to the IETF RFC 9000 draft.
- Added support for opaque symmetric key objects in the form of the EVP_SKEY data structure.
- Disabled the SHA-224 digest.
- SHAKE-128 and SHAKE-256 implementations no longer have a default digest length. Therefore, these algorithms cannot be used with the **EVP_DigestFinal/_ex()** function unless the **xoflen** parameter is set.
- Added a capability for a client to send multiple key shares in TLS 1.3 connections.

Jira:RHEL-80854[1]

OpenSSL supports sslkeylogfile

OpenSSL supports the **sslkeylogfile** format for TLS. As a result, you can log all secrets produced by SSL connections by setting the **SSLKEYLOGFILE** environment variable.

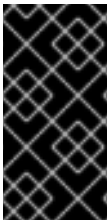

IMPORTANT

Enabling the **SSLKEYLOGFILE** variable poses an explicit security risk. Recording the exchanged keys during an SSL session allows anyone with read access to the file to decrypt application traffic sent over that session. Use this feature only in test and debug environments.

Jira:RHEL-90854

Hybrid ML-KEM cryptography works in FIPS mode

With this release, Hybrid Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) post-quantum cryptographic algorithms are supported in FIPS mode of RHEL. OpenSSL is able to fetch the Elliptic Curve Diffie-Hellman (ECDH) part of the new hybrid post-quantum groups from the FIPS provider when the system is running in FIPS mode. As a result, the OpenSSL library uses FIPS-compliant cryptography for the ECDH part of the hybrid post-quantum key exchanges. When you set the system to the **FIPS:PQ** cryptographic policy, the hybrid post-quantum groups are enabled and used by default by OpenSSL servers and clients.

Jira:RHEL-95239

crypto-policies support Ed25519 in NSS

With this update to the system-wide cryptographic policies, support for the SHA-512 variant of the Edwards-curve Digital Signature Algorithm (EdDSA), Ed25519, is available for Network Security Services (NSS). As a result, **crypto-policies** enable Ed25519 in DEFAULT, LEGACY, and FUTURE policies for NSS by default.

Jira:RHEL-104607

New package: **rust-rpm-sequoia**

RHEL 9.7 introduces the **rust-rpm-sequoia** package to support quantum-resistant signatures in RPM packages through the **multisig** DNF plug-in. This addition enables you to verify OpenPGP v6 signatures in RPM packages signed with post-quantum cryptographic (PQC) algorithms.

Jira:RHEL-126412[1]

SCAP Security Guide rebased to 0.1.78

For additional information, see the SCAP Security Guide release notes .

Jira:RHEL-111009

The SELinux policy adds rules and type for the **qgs** daemon

The **qgs** daemon was added to RHEL with the **linux-sgx** package, which supports TDX confidential virtualization. The **qgs** daemon communicates with QEMU over a UNIX domain socket when the guest OS requests attestation of the virtual machine (VM). To make this possible, the SELinux policy adds a new **qgs_t** type, access rules, and permissions.

Jira:RHEL-87744

Three RHEL services removed from SELinux permissive mode

The following SELinux domains for RHEL services have been removed from SELinux permissive mode:

- **powerprofiles_t**

- **samba_bgqd_t**
- **switcheroo_control_t**

Previously, these services from packages recently added to RHEL 10 were temporarily set to SELinux permissive mode, which allows gathering information about additional denials while the rest of the system is in SELinux enforcing mode. This temporary setting has now been removed, and as a result, these services now run in SELinux enforcing mode.

Jira:RHEL-82674[1]

tuned-ppd confined in the SELinux policy

RHEL 9.7 adds additional rules to the SELinux policy that confine the **tuned-ppd** service. Before this update, the service ran with the **unconfined_service_t** SELinux label, which violated the CIS Server Level 2 benchmark "Ensure No Daemons are Unconfined by SELinux" rule. With this update, the service is no longer unconfined and runs successfully in SELinux enforcing mode.

Jira:RHEL-69526

Keylime rebased to version 7.12.1

The Keylime packages have been rebased to upstream version 7.12.1. The most important fixes and enhancements include:

- Implemented security fix for CVE-2025-1057 addressing vulnerability of the registrar component when updated to version 7.12.0.
- Added support for named measured boot policies, which makes policy organization easier.
- Fixed resource handling in webhook operations.
- Fixed certificate generation to follow the X.509 v3 certificate and X.509 v2 certificate revocation list (CRL) standards according to RFC 5280.

Jira:RHEL-78418

SELinux assigns a particular type to /dev/diag

With this update, the **diagnostic_device_t** type is assigned to the **/dev/diag** device in the SELinux policy. As a result, SELinux can properly control access to the device.

Jira:RHEL-95342[1]

OpenSSL PKCS #11 provider adds support for Ex=RSA ciphers

This update of the OpenSSL PKCS #11 provider enables the use of PKCS #11 tokens with OpenSSL without relying on deprecated functionality. This alternative resolves the unsupported RSA padding mode issue, ensuring seamless use of Ex=RSA ciphers with hardware security modules (HSMs) on RHEL 9. This results in eliminating TLS handshake failures and providing secure communication when establishing TLS 1.2 connections with OpenSSL and PKCS #11 tokens.

Jira:RHEL-105625

New package: fips-provider-next

The **fips-provider-next** package provides the next version of the FIPS provider that is submitted to the National Institute of Standards and Technology (NIST) for validation. The package is not installed by

default because the **openssl-fips-provider** package is the validated OpenSSL FIPS provider. To switch from **openssl-fips-provider** to **fips-provider-next**:

```
# dnf swap openssl-fips-provider fips-provider-next
```

Jira:RHEL-105009

Rsyslog imuxsock provides the new ratelimit.discarded counter

With this update, the **imuxsock** Rsyslog module includes a new counter, **ratelimit.discarded**, which tracks the number of messages dropped due to rate-limiting on the Unix socket. This enhancement provides administrators with visibility into message loss due to rate-limiting, enabling them to fine-tune their rate-limiting settings and prevent critical logs from being discarded.

Jira:RHEL-66274

Rsyslog imfile provides the new deleteStateOnFileMove option

With this update, the new **deleteStateOnFileMove** parameter has been added to the **imfile** module, available as both a module-level and a per-action option. This enhancement addresses the issue of orphaned state files accumulating in the **spool/** directory when monitored log files are rotated or moved. By enabling this parameter, you can automatically clean up these obsolete files when log files are moved, preventing disk space from being wasted and simplifying management.

Jira:RHEL-92262[1]

4.3. SUBSCRIPTION MANAGEMENT

Simplified status for systems registered to SCA-enabled organizations

Before this update, when registered to a Simple Content Access (SCA) enabled organization, the **subscription-manager status** command reported **Overall Status: Disabled** and **System Purpose Status: Disabled**. Because this status was confusing and often misinterpreted as an error, the status report has been simplified. Now the **Overall Status** reports either **Registered** or **Not registered** and **System Purpose Status** has been eliminated.

For more information on SCA, see Simple Content Access .

Jira:RHEL-84890[1]

4.4. SOFTWARE MANAGEMENT

dnf4 can be used to run DNF commands

With this update, you can enter either **dnf** or **dnf4** to run DNF commands.

Jira:RHEL-82310

DNF can verify RPMv6 signatures on RPM packages

Quantum-safe cryptography guarantees integrity and origin of software. However, in quantum computing, standard asymmetric cryptography algorithms, such as RSA, are no longer relevant. With this update, you can use the new **multisig** DNF plugin to verify RPMv6 signatures on RPM packages, in addition to standard RPMv4 signatures. RPMv6 signatures can be based on quantum-safe algorithms, such as ML-DSA.

To verify RPMv6 signatures, you can install the **multisig** plugin through the **python3-dnf-plugin-multisig** RPM package.

NOTE

Successful verification is a prerequisite for installing, reinstalling, upgrading, or downgrading packages from a repository that has the **gpgcheck** option set to **True**.

Jira:RHEL-100157

createrepo_c supports zstd

This enhancement adds support for the **Zstandard** (**zstd**) compression algorithm for **createrepo_c** commands. As a result, **createrepo_c** can read and generate metadata compressed with **zstd**.

Jira:RHEL-67689

dnf marks transient transactions in DNF history

The **dnf history info** command shows whether a transaction was persistent or transient. As a result, it is easier to keep track of package changes, especially on systems with many transient packages.

Jira:RHEL-84512

RPM records a checksum of the original package during installation

With this update, RPM records the SHA256 and SHA512 digests of the entire **.rpm** package during its installation. You can then retrieve these digests from the RPM database to verify that the installed package corresponds to a specific **.rpm** file. As a result, you can improve the integrity of your RHEL system by retrospectively verifying that the installed package set matches, bit-by-bit, a known set of **.rpm** packages, such as the ones available in a DNF repository.

To print the package digests of an installed package, use the following command:

```
$ rpm -q --qf "[%{packagedigestalgos:hashalgo} %{packagedigests}\n]" <package_name>
```

You can also customize which digest types are recorded in the database by configuring the new **%_pkgverify_digests** macro, for example:

```
%_pkgverify_digests 8:10
```

Jira:RHEL-35619

RPM supports spec-local file attributes and dependency generators

File attributes and their dependency generators are usually shipped in separate packages that you must install prior to building a package that uses these attributes. However, you might need a file attribute to take effect during the build of the package that ships this attribute. You might also need the file attribute just for building the package, without shipping the attribute at all.

With this update, you can register **spec**-local file attributes and generators by performing the following actions:

1. Define the **%_local_file_attrs** macro. **%_local_file_attrs** accepts a colon-separated list of new attribute names to register directly in your **spec** file.

2. Define one or more dependency generator macros for each attribute, such as **%__NAME_provides** or **%__NAME_path**, where **NAME** is the name of the local file attribute.

RPM then uses the file attributes for dependency generation when the **spec** file is built. As a result, you can create build-time file attributes that are not necessarily meant for installation.

For example, the following **spec** file snippet generates the provides for each packaged file by using the **foobar.sh** script bundled with your package's sources:

```
Source1: foobar.sh
[...]
%define _local_file_attrs foobar
%define __foobar_provides %{SOURCE1}
%define __foobar_path .*
```

Jira:RHEL-52772

New $releasever_major and $releasever_minor variables

The new **$releasever_major** and **$releasever_minor** variables are available to better support the Extra Packages for Enterprise Linux (EPEL) repository and other repositories that distribute content per major version of RHEL instead of per minor version. These variables are automatically derived from the **$releasever** variable or the **system-release(releasever_major)** and **system-release(releasever_minor)** virtual provides. As a result, you can use **$releasever_major** and **$releasever_minor** to create repository configuration files that work across multiple major or minor versions of RHEL.

Jira:RHEL-65817

4.5. SHELLS AND COMMAND-LINE TOOLS

openCryptoki provided in version 3.25.0

The **openCryptoki** packages are provided in version 3.25.0. Support has been added for the following:

- In EP11:
  - PKCS#11 v3.0 SHA3 and SHA3-HMAC mechanisms
  - PKCS#11 v3.0 SHA3 mechanisms and MGFs for RSA-OAEP
  - PKCS#11 v3.0 SHA3 variants of RSA-PKCS and ECDSA mechanisms
  - Opaque secure key blob import via C_CreateObject
- In ICA/Soft:
  - PKCS#11 v3.0 SHAKE key derivation
  - The CKM_AES_KEY_WRAP[_*] mechanisms
  - The CKM_ECDH_AES_KEY_WRAP mechanism
  - Key wrapping with AES-GCM
- In CCA:

- CCA AES CIPHER secure key types
- The CKM_ECDH1_DERIVE mechanism
- Newer CCA versions on s390x and non-s390x platforms
- CKM_AES_GCM for single-part operations only

- CCA/Soft/ICA: The CKM_RSA_AES_KEY_WRAP mechanism.
- P11KMIP: Added a tool for importing and exporting PKCS#11 keys to a KMIP server.
- ICA: Report mechanisms depending on whether libica is in FIPS mode.

Jira:RHEL-73344[1]

GIMP rebased to 3.0.4

The GNU Image Manipulation Program (GIMP) has been rebased to stable upstream version 3.0.4 in RHEL 9.7.

Jira:RHEL-40106[1]

4.6. INFRASTRUCTURE SERVICES

RHEL is now equipped with **dyninst** version 13.0.0

The **dyninst** framework is rebased to upstream version 13.0.0 This version offers the following list of enhancements:

- improved support for AMD GPU binaries.
- improved parsing of x86 instructions and C++ DWARF constructs.

For more information, see the upstream documentation.

Jira:RHEL-87002

RHEL is now equipped with **SystemTap** version 5.3

SystemTap is rebased to version 5.3, and its multithreaded parsing capability now improves startup performance by reducing initialization time by several seconds.

Jira:RHEL-87000

elfutils is now rebased to version 0.193

elfutils 0.193 is now available in RHEL 9.7. The notable changes in this update include:

- **debuginfod** now supports CORS (webapp access) in the web API and provides a **--cors** option. The new **--listen-address** option enables binding the HTTP listen socket to a specific IPv4 or IPv6 address. The **debuginfod** client now caches **x-debuginfod-*** HTTP headers alongside downloaded files.
- **libdw** library adds the **dwarf_language** and **dwarf_language_lower_bound** functions, with improved support for DWARF6 language metadata and new language constants for Nim, Dylan, Algol68, V, and Mojo. The **dwarf_srclang** function is forward-compatible with DWARF6

language constants.

- **libdwfl_stacktrace** experimental interface can unwind stack samples into call chains and cache ELF data for multiple processes. This interface initially supports **perf_events** stack sample data and is provided as a Technology Preview.
- **libelf** library has a more robust implementation of **elf_scnshndx** for ELF files with more than 64K sections.
- **readelf** tool improves handling of corrupt ELF data. The output of the **--section-headers** option now includes a key to explain section flag meanings.

Jira:RHEL-86971

valgrind has been upgraded to upstream version 3.25.1

The upgrade from version 3.24.0 (RHEL 9.6) to the upstream version 3.25.1 (RHEL 9.7) provides the following notable enhancements:

- Added support for zstd-compressed debug sections.
- Extended support to Linux syscalls: landlock*, io_pgetevents, open_tree, move_mount, fsopen, fsconfig, fsmount, fspick, userfaultfd.
- Enhanced file-descriptor tracking: **--track-fds=yes** and **--track-fds=all** apply the same behavior to inherited file descriptors as to standard input, standard output, and standard error.
- New option **--modify-fds=high** (use with **--track-fds=yes**) allocates higher-numbered descriptors first to help detect descriptor reuse issues.
- Helgrind configuration: warnings for **pthread_cond_signal** and **pthread_cond_broadcast** with an unlocked mutex are now controlled by **--check-cond-signal-mutex=yes|no** (default: no).

Architecture-specific enhancements:

- New IBM Z (**s390x**) NNPA hardware support.

Jira:RHEL-86998

valgrind package split into subpackages for flexible installation

Before this update, the **valgrind** package included all components in a single package. As a consequence, you had to install features that you did not need.

With this update, the **valgrind** package has been split into multiple subpackages. As a result, you can install only the required components you require, such as the core **valgrind** functionality, postprocessing scripts, GDB integration, or documentation.

Jira:RHEL-75468[1]

Valkey 8 is now available

Valkey 8, an advanced key-value store, is now available in RHEL. It functions as a data structure server, allowing keys to store various data types, for example:

- Strings

- Hashes
- Lists
- Sets
- Sorted sets

Valkey is fully compatible with clients and serves as an alternative to Redis.

Jira:RHEL-89978[1]

fs.protected_regular and fs.protected_fifos sysctls parameters are enabled by default

Previously, in the RHEL 9 kernel the **fs.protected_regular** and **fs.protected_fifos** sysctls parameters were added to make some data spoofing attacks harder. Now, these parameters are enabled by default which improves the security for installations. To disable these **sysctls** parameters, add the following lines in the **/etc/sysctl.d/60-protected.conf** file:

- **fs.protected_regular = 0**
- **fs.protected_fifos = 0**

Jira:RHEL-50534[1]

The BrowseOptionsUpdate directive is now available in RHEL

The **BrowseOptionsUpdate** directive determines the source and update frequency of default printing options. It specifies whether the system retrieves options from a local system or a remote printing server, and if it updates them at service startup, at certain intervals, or not at all.

You can now add the **BrowseOptionsInterval** directive and its value to the **/etc/cups/cups-browsed.conf** file to achieve the required behavior. The directive offers these values:

- **None** (default): A local file, created from previous sessions, loads default options.
- **Static**: The **cups-browsed** service retrieves default options from the remote server when it starts.
- **Dynamic**: The system updates default options according to the **BrowseInterval** value, also defined in the **/etc/cups/cups-browsed.conf** file.

Note: You need to restart the service after changing the **BrowseOptionsInterval** directive values.

Jira:RHEL-6519[1]

RHEL 10 provides gpsd in version 3.26.1

In RHEL 10, the **gpsd tools** package is provided in version 3.26.1. This version offers improved support for u-blox receivers.

Jira:RHEL-90132[1]

4.7. NETWORKING

Nmstate can assign settings to network interfaces based on PCI addresses

With this enhancement, you can use Nmstate to set up network interfaces based on their PCI address instead of a device name. Use this feature to ensure consistent configuration across nodes in a cluster. For further details, see Configuring an Ethernet connection with a dynamic IP address by using nmstatectl with a device path and Configuring an Ethernet connection with a static IP address by using nmstatectl with a device path.

Jira:RHEL-88993

Bond configurations in Nmstate support optimization settings

With this enhancement, the Nmstate API supports the following bond options:

- **lacp_active**: Defines whether or not the Linux kernel periodically sends Link Aggregation Control Protocol Data Unit (LACPDU) frames. You can use this setting only in the 802.3ad bond mode.
- **ns_ip6_target**: Lists the IPv6 addresses to use as IPv6 monitoring peers when you set the **arp_interval** parameter to a value larger than 0.

As a result, administrators can use these settings to optimize a network bond to ensure stable connections, efficient bandwidth, and IPv6 compatibility.

Jira:RHEL-85784

nmtui now supports configuring the loopback interface

NetworkManager already supports configuring the loopback interface by using the **nmcli** utility. This enhancement adds the same functionality to the **nmtui** application. As a result, you can configure IP addresses and routes on the loopback interface.

Jira:RHEL-85770

The NetworkManager-libreswan plugin supports using the Libreswan default values

With this enhancement, you can set the **no-nm-default** property in Libreswan VPN connection profiles to **true** to use Libreswan's instead of NetworkManager's default values. This ensures the compatibility with configurations defined for native Libreswan. As a result, you can now, for example, configure subnet-to-subnet tunnels.

Jira:RHEL-85768

NetworkManager now supports fixed subnet IDs for downstream interfaces when using IPv6 prefix delegation

With this enhancement, you can now specify a fixed subnet ID for downstream interfaces in NetworkManager when you use IPv6 prefix delegation. In previous releases, when you rebooted the system, the subnet ID for these interfaces could change. With a fixed subnet ID, IPv6 addresses assigned to devices in the downstream network do not change when you reboot the RHEL host.

Jira:RHEL-85765

An NBFT parser was added to nm-initrd-generator

NVMe Boot Firmware Table (NBFT) is a standard method for firmware to pass network and storage configuration from the pre-boot environment directly to the operating system by using an ACPI table. The **nm-initrd-generator** utility now uses this parser to automatically detect and apply this

configuration, and creates the necessary connections without manual setup. This implementation replaces the **95nvmf** module in **dracut** and relies on **systemd** automation for a more streamlined and robust boot sequence.

Jira:RHEL-83061

Nmstate now supports configuring FEC settings for network interfaces

With this enhancement, you can now use Nmstate to apply Forward Error Correction (FEC) modes, such as **RS-FEC**, **Base-R** and **Disabled** to interfaces. These settings are crucial for improving data transmission reliability by detecting and correcting errors without retransmission. As a result, you can now use Nmstate to apply FEC settings instead of manually configuring them or using platform-specific tools.

Jira:RHEL-80725[1]

Nmstate now supports the mtu and quickack route options

With this enhancement, you can use Nmstate to set the **mtu** and **quickack** route options. These settings are important for optimizing the network performance if the maximum transmission unit is different from the default and for tuning the TCP acknowledgment behavior. As a result, you now have more precise control over network traffic behavior.

Jira:RHEL-80418

The mlx5 driver now supports symmetric OR-XOR RSS hash

With this enhancement, the default transform (**xfrm**) for Receive Side Scaling (RSS) is now **symmetric-or-xor**.

Due to this new default, modifying the **rx-flow-hash** setting by using the **ethtool** utility now requires one of the following actions:

- Set **rx-flow-hash** to a value that is compatible with symmetric hashing: **sdfn**, **sd**, or **fn**.
- Set **xfrm** to **none** before setting a different **rx-flow-hash** value, for example:

```
# ethtool -X enp0s1 xfrm none
# ethtool -N enp0s1 rx-flow-hash udp4 n
```

Jira:RHEL-73517[1]

ModemManager rebased to version 1.22

The **ModemManager** packages have been upgraded to upstream version 1.22. This version includes bug fixes and support for additional devices.

For a complete list of changes, see the upstream release notes.

Jira:RHEL-68732

Nmstate now supports egress and ingress priority mapping for VLAN interfaces

NetworkManager already supports configuring traffic priority mapping for VLAN interfaces. With this enhancement, the Nmstate library can also handle both egress and ingress priority quality of service (QoS) mapping rules. As a result, you can use Nmstate to create VLANs and define bidirectional priority mapping, helping manage traffic more precisely and efficiently.

Jira:RHEL-67631

Nmstate now supports configuring routes by using a MAC address instead of an interface name

With Nmstate, you can create a network connection by assigning it to the MAC address of an interface. With this enhancement, you can use the profile name instead of the interface name in the **next-hop-interface** parameter in the routing configuration. With this feature, you can create static routes without knowing the interface name.

Jira:RHEL-32495

New network packet drop reasons and MIB counters

The kernel's networking stack now provides more detailed reasons when it drops network packets. This enhancement also adds two new Management Information Base (MIB) counters: **LINUX_MIB_PAWS_TW_REJECTED** and **LINUX_MIB_PAWS_OLD_ACK**. As a result, debugging and diagnosing network problems, is now easier.

Jira:RHEL-88890[1]

The fwctl subsystem has been added to the kernel

If the kernel lock-down feature is enabled, the kernel does not allow access to **resource0** files in the /**sys**/ directory and PCI config spaces for security reasons. The **fwctl** kernel subsystem manages communication with the firmware in software-defined devices, such as the **mlx5** network interface controller. This subsystem establishes a standardized and secure Remote Procedure Call (RPC) interface, that enables user-space applications to interact with device firmware for diagnostics, configuration, and updates. In addition to the new subsystem, the **mstflint** utility now also uses the **fwctl** subsystem, and the utility functions fully in these secure environments.

Jira:RHEL-86016[1]

The ice driver now supports reducing the MSI-X vector usage for a PF to free vectors for associated VF

With this enhancement, you can now reduce the Message Signaled Interrupts eXtended (MSI-X) vectors allocated to a physical function (PF) to ensure that a sufficient number of vectors are available for associated virtual functions (VFs). For details, see Reducing the MSI-X vector usage for a physical function to free vectors for associated virtual functions.

Jira:RHEL-63642[1]

iproute rebased to version 6.14.0

The **iproute** package has been updated to upstream version 6.14.0.

Notable enhancements:

- The **ip nexthop** command supports 16-bit **nexthop** weights.
- The **ip link rmnet** command supports flag handling.
- The **ip lwtunnel** command supports setting and getting the 'tunsrc' attribute.
- The **ip monitor** command adds support for monitoring multicast addresses (**ip monitor maddress**).

- The **ip rule** command supports the 'dscp' selector.
- The **ip rule** command supports flow labels.
- The **ip route** command supports IPv6 flow labels.
- The **ip address** and **ip link show** commands support the 'down' filter.
- The **tc flower** filter supports matching on tunnel metadata.
- The **tc fq** queuing discipline supports the **TCA_FQ_OFFLOAD_HORIZON** attribute.
- The **tc** utility supports the **Hold**/**Release** mechanism in Time-Sensitive Networking (TSN) as specified in the IEEE 802.1Q-2018 standard.
- The **rdma monitor** command adds support for monitoring Remote Direct Memory Access (RDMA) events.
- The **vdpa** utility supports setting the MAC address.
- Several man pages were improved.

Notable bug fixes:

- Some memory leaks were fixed.
- The error checking of the **ip netconf** command was fixed to prevent unnecessarily strict errors.
- Custom **iproute2** settings in the **/etc/iproute2/** directory work as expected.

Jira:RHEL-90492

4.8. KERNEL

Kernel version in RHEL 9.7

Red Hat Enterprise Linux 9.7 is distributed with the kernel version 5.14.0-611.5.1.

Added support for virtio devices

Before this update, **virtio** devices inside of KVM guests were all listed as type **generic-ccw**. With this enhancement, you can easily identify which device type is connected at which device number by using the **lszdev** command:

```
# lszdev
TYPE ID ON PERS NAMES

virtio-balloon 0.0.0007 yes no
virtio-blk 0.0.0000 yes no vda
virtio-console 0.0.0004 yes no
virtio-gpu 0.0.0002 yes no
virtio-input 0.0.0005 yes no
virtio-input 0.0.0006 yes no
virtio-net 0.0.0001 yes no enc1
virtio-scsi 0.0.0003 yes no
virtio-vsock 0.0.0008 yes no
```

This enhancement also introduces additional **chpstat** fixes for Red Hat Enterprise Linux 9.4 and 9.6, improving DPU utilization scaling in reports (**s390utils** and **s390-tools**).

Jira:RHEL-73342[1]

kpatch-dnf plugin is updated with improved kernel management

Before this update, the **kpatch-dnf** plugin did not align kernel upgrades with **kpatch** support. As a consequence, administrators might install or upgrade to kernels that were not supported by **kpatch**, thereby increasing the risk of running unsupported kernels and reducing system stability.

With this update, the **kpatch-dnf** plugin enables administrators to focus kernel updates on those supported by **kpatch**. As a result, system upgrades are more reliable, and overall stability is improved.

Jira:RHEL-85579[1]

Arm SPE support extended to Neoverse-V2 and Cortex CPUs in the kernel

The Arm SPE feature support in **kernel** has been extended to include Neoverse-V2 and Cortex CPUs. As a result, users can now access Arm SPE capabilities for improved observability and analysis when running workloads on Neoverse-V2 and Cortex CPUs.

Jira:RHEL-60216[1]

Intel Arrow Lake U RAPL energy events support in kernel

Before this update, the Intel Arrow Lake U microarchitecture did not support RAPL (Running Average Power Limit) energy performance counters in the **kernel** package. As a result, users could not monitor or measure energy consumption for Arrow Lake U systems using standard perf tooling.

With this update, support for RAPL energy events is added for Arrow Lake U in the **kernel** package. The perf tool identifies power consumption events for Arrow Lake U platforms. You can now monitor energy usage for CPU cores, GPUs, packages, and system domains.

Jira:RHEL-53585[1]

Added support for core energy counters in kernel

The kernel supports per-core energy measurement on AMD CPUs. The Power Management Unit (PMU) exposes the **power_core** PMU and the **energy-core** event so that you can monitor energy consumption for each CPU core. This enhancement aligns with AMD per-core energy counter capabilities.

Jira:RHEL-52654[1]

Perf support for Intel Clearwater Forest core counters

Before this update, you could not monitor hardware events on Intel Clearwater Forest CPUs by using **perf** core counters. With this update, the **perf** package recognizes the Clearwater Forest Performance Monitoring Unit (PMU). It provides named core events, including Topdown Level 1 metrics, such as front-end bound, back-end bound, retiring, and slots. Perf also uses architectural process event-based sampling (PEBS) on this microarchitecture to provide low-overhead sampling of selected events. As a result, you can collect core counter data and perform Top-down analyses on Clearwater Forest systems.

Jira:RHEL-47454[1]

Adaptive PEBS enables counter snapshotting support in perf on Intel Panther Lake

Before this update, the Linux kernel's **perf** tool relied on software-based sample reads to collect performance event data, which introduced minor timing gaps and additional overhead when reading counters after an event overflow. With this update, adaptive PEBS counter snapshotting is available on Intel Panther Lake CPUs. This hardware feature enables the kernel to capture programmable counters, fixed-function counters, and performance metrics directly in the PEBS record using the PEBS format version 6.

As a result, counter snapshotting provides a more accurate and lower-overhead alternative to software sample reads, improving performance monitoring and analysis capabilities.

Jira:RHEL-47444[1]

Intel Trace Hub supports Intel Panther Lake

This update adds Intel Trace Hub device IDs for the Panther Lake platforms (P, H, and U). The systems based on Panther Lake can use Intel Trace Hub features for debugging and tracing.

Jira:RHEL-47424[1]

Perf uncore event support for Intel Clearwater Forest

Before this update, uncore event monitoring was not available for Intel Clearwater Forest microarchitecture. With this update, the **perf** package supports uncore event monitoring on Clearwater Forest systems. As a result, you can perform advanced performance analysis and debugging on supported hardware.

Jira:RHEL-45095[1]

Intel Arrow Lake H microarchitecture support added to intel_th

Before this update, Intel Trace Hub did not recognize Arrow Lake H NPK device IDs, which limited trace and debugging capabilities for systems that use this hardware. With this update, the **intel_th** package supports the Intel Arrow Lake H microarchitecture in Intel Trace Hub. As a result, you have enhanced tracing and debugging features on Arrow Lake H platforms.

Jira:RHEL-20110[1]

PerfMon support enabled for Intel Arrow Lake H in kernel

With this update, the **kernel** package provides PerfMon support for Core, Uncore, Cstate, and MSR features on the Intel Arrow Lake H microarchitecture. As a result, you can monitor and analyze performance metrics specific to Arrow Lake H systems by using the **perf** tool.

Jira:RHEL-20094[1]

Enhanced pstore functionality in virtual and cloud environments

The **pstore** kernel feature, which saves crash and panic information persistently, is now easier to use in virtualized environments and cloud platforms. With this release, you can enable the use of EFI variables for **pstore** without the **efi_pstore.pstore_disable=0** kernel parameter while the system is running:

```
$ echo "N" > /sys/module/efi_pstore/parameters/pstore_disable
```

This enhancement simplifies the activation and post-crash data retrieval for **pstore**, improving troubleshooting and system reliability in environments where the ACPI ERST method is unavailable.

Jira:RHEL-2564[1]

The default measurement module for rteval is now rtla timerlat for better tracing of problem latencies

With this enhancement, you should be able to easily identify the source of problem latencies. The desired cyclictest measurement module can be chosen using the rteval.config file.

Jira:RHEL-97540[1]

KVM modules are integrated into the Realtime Kernel package

This update removes the generation of KVM module packages for the Realtime Kernel in RHEL, aligning with the decision to make the Realtime Kernel a deployment option for base RHEL. This change streamlines the deployment process, integrating KVM modules directly into the Realtime Kernel package and eliminating the separate **kernel-rt-kvm** package. As a result, users will experience a more seamless and efficient setup when deploying the Realtime Kernel on RHEL, improving the overall user experience.

Jira:RHEL-76757[1]

kernel supports Shadow Stack (SHSTK) Ring 3 kernel

Before this update, the **kernel** package did not support Shadow Stack (SHSTK) in Ring 3 for **x86_64** architectures. As a consequence, user-space applications could be vulnerable to control flow hijacking attacks.

With this update, the **kernel** package introduces Control-flow Enforcement Technology (CET) Shadow Stack support for Ring 3. This enhancement provides a hardware-enforced secondary stack that cannot be directly modified by applications. As a result, applications running on supported Intel Sapphire Rapids processors now have improved protection against control flow attacks in the user space.

Jira:RHEL-15599[1]

python-drgn rebased to version 0.0.31

python-drgn has been rebased to version 0.0.31. This update introduces several enhancements and new features:

- Added support for **debuginfod**, which enables automatic retrieval of debugging information from debuginfod servers.
- A new Module API, which provides improved extensibility and integration capabilities.
- Kernel stack unwinding without debugging symbols, allowing stack traces to be generated even when debug symbols are unavailable.

For a complete list of changes, see the upstream changelogs:

- 0.0.31: https://github.com/osandov/drgn/releases/tag/v0.0.31
- 0.0.30: https://github.com/osandov/drgn/releases/tag/v0.0.30

Jira:RHEL-86264

crash rebased to 9.0.0

The **crash** package, which provides a kernel analysis utility for live systems and various types of dump files, has been rebased to upstream version 9.0.0. This version provides a number of fixes and enhancements, most notably the following:

- The internal **gdb** database has been updated to version 16.2.
- The **crash** utility now supports cross-compilations.

Jira:RHEL-76270

Support for per-core energy tracking (RAPL perf events) for AMD CPUs

With this enhancement, the addition of the core RAPL counter support is added. As a result, the AMD systems can measure the core-level power information in addition to the package-level power information.

Jira:RHEL-23496[1]

Default configuration now disables jitter entropy source in rng-tools

The jitter entropy source is now disabled by default in **rng-tools**. Modern CPUs typically provide a hardware entropy source, and most virtual machines offer the **/dev/hwrng** device as an entropy source from the virtual host. In these environments, the jitter entropy source consumes unnecessary CPU cycles. For older hardware without a hardware entropy source, you can explicitly enable the jitter entropy source in **/etc/sysconfig/rngd**.

As a result, the **rngd** daemon no longer consumes CPU cycles unnecessarily on systems that have hardware entropy sources.

Jira:RHEL-91119

NVMf-FC kdump is now supported on the IBM Power

NVMf-FC kdump now supports the IBM Power system for running **kexec-tools**. This allows the capture of system memory dumps over a fiber channel network using the NVMe storage devices for high-speed and low-latency access to storage for crash dump data.

Jira:RHEL-11471[1]

4.9. BOOT LOADER

Secure boot on aarch64 enabled through Microsoft-signed shim

The **shim** package for the 64-bit ARM architecture is signed by Microsoft to enable secure boot by default on platforms that trust the Microsoft UEFI CA. This aligns the ARM boot path with x86 and removes the need to add custom **PK**, **KEK**, or **db** entries.

Before this update, RHEL 9 on the 64-bit ARM architecture could not use secure boot on cloud and vendor platforms that rely on Microsoft's UEFI trust chain. This blocked standard, compliant deployments, including on Google Compute Engine.

Starting from RHEL 9.7, secure boot works by default on RHEL 9 for the 64-bit ARM architecture. Direct and fallback boot paths are successful, and the associated EFI binaries are correctly signed.

Jira:RHEL-18969[1]

4.10. FILE SYSTEMS AND STORAGE

multipathd supports file-based sockets

With this update, the **multipathd** daemon listens for commands on a file-based socket **/run/multipathd.socket** in addition to the abstract namespace socket. You can communicate with the host's **multipathd** daemon from within a container by using a bind mount for the new socket file.

Jira:RHEL-78758[1]

Automatic RAID checks are enabled by default

With this update, the **raid-check** service is enabled by default. This ensures that **raid-check.service** runs automatically at scheduled intervals after the system boots, performing periodic RAID consistency checks without requiring manual intervention.

Jira:RHEL-86164

LVM RAID repairs volumes after multiple simultaneous device failures

With this enhancement, you can use the **lvconvert --repair /dev/*VG-name/LV-name*** command to reintegrate missing RAID devices back into a striped RAID (raid4, raid5, and raid6). This repair process works even when the number of temporarily missing devices exceeds the fault tolerance of the RAID level, allowing for recovery once the devices reappear. Note that you must unmount and deactivate the volume and the file system on top before repairing them.

Jira:RHEL-67039

4.11. HIGH AVAILABILITY AND CLUSTERS

New resource agent for managing etcd in Podman containers is available

Before this update, Red Hat High Availability did not provide a resource agent for managing **etcd** running in Podman containers.

With this enhancement, the new **podman-etcd** resource agent has been added.

As a result, you can create and manage resources for **etcd** running in a Podman container. This agent is a required component for the Two Node OpenShift with Fencing (TNF) solution.

Jira:RHEL-88429

The Filesystem resource agent supports the aznfs file system type

Before this update, to manage an Azure Network File System file share in a cluster, you had to configure the **Filesystem** resource agent with **fstype=nfs**. This method did not support Azure-specific features, such as Encryption in Transit.

With this update, the **Filesystem** resource agent supports **aznfs** as a file system type.

As a result, you can set **fstype=aznfs** when creating a **Filesystem** resource to manage an Azure Network File System file share. This enables support for Azure-specific features. Note that this functionality requires the **aznfs** client package from the Microsoft repository to be installed on all cluster nodes.

Jira:RHEL-88035

Oracle Database 23ai is supported as a cluster resource

Before this update, the Oracle database resource agent was not tested for use with the Oracle Database 23ai release. Therefore, this version was not supported as a highly available resource within a Pacemaker cluster.

With this update, the existing Oracle resource agent has been successfully tested and validated with Oracle Database 23ai.

As a result, Pacemaker supports managing Oracle Database 23ai instances, enabling fully tested high availability configurations for this version.

Jira:RHEL-85220[1]

The fence_sbd agent can automatically detect the SBD device

Before this update, when configuring a **fence_sbd** resource, you were required to explicitly specify the SBD device path by using the **devices** parameter.

With this update, the **fence_sbd** agent can now retrieve the device configuration from the system.

As a result, if you do not set the **devices** parameter when creating the **fence_sbd** resource, the agent automatically uses the device specified in the **SBD_DEVICE** variable within the **/etc/sysconfig/sbd** file.

Jira:RHEL-79798

Watchdog device listing provides more detailed information

Before this update, when listing available watchdog devices, the output only displayed the device path, such as **/dev/watchdog0**. This made it difficult for administrators to distinguish between multiple devices on the same system.

With this update, the output includes the device path, identity, and driver for each watchdog. This allows for easy identification and selection of the correct device.

Jira:RHEL-76177

pcs warns users before removing the last fencing device

Before this update, **pcs** allowed users to disable or remove the last fencing device from a cluster without a warning. This could inadvertently leave the cluster in an unsupported state without any STONITH or SBD fencing configured.

With this enhancement, **pcs** now includes a safety check to prevent the accidental removal of all fencing mechanisms.

As a result, if you attempt an action that would leave the cluster without any fencing, **pcs** displays an error and blocks the change by default. For example, this occurs when you try to remove the last STONITH resource while SBD is disabled. You can override this safety check to force the change if needed.

Jira:RHEL-76170

The pcs node attribute and pcs node utilization commands now support multiple output formats

Previously, the **pcs node attribute** and **pcs node utilization** commands displayed their output only in a human-readable plain text format. This format was not suitable for machine parsing or for easily replicating the configuration.

With this enhancement, a new **--output-format** option has been added to the **pcs node attribute** and **pcs node utilization** commands.

As a result, you can now display the configured node attributes and utilization in one of three formats:

- **text**: Displays the output in plain text. This is the default format.
- **json**: Displays the output in a machine-readable JSON format, which is useful for scripting and automation.
- **cmd**: Displays the output as a series of **pcs** commands, which you can use to recreate the same configuration on a different system.

Jira:RHEL-76154

The pcs alert config command now supports multiple output formats

Previously, the **pcs alert config** command displayed its output only in a human-readable plain text format. This format was not suitable for machine parsing or for easily replicating the configuration.

With this enhancement, a new **--output-format** option has been added to the **pcs alert config** command.

As a result, you can now display the configured alerts in one of three formats:

- **text**: Displays the output in plain text. This is the default format.
- **json**: Displays the output in a machine-readable JSON format, which is useful for scripting and automation.
- **cmd**: Displays the output as a series of **pcs** commands, which you can use to recreate the same alert configuration on a different system.

Jira:RHEL-76153

pcs automatically validates the CIB for potential issues

Previously, the **pcs** utility did not automatically run advanced validation checks on the Cluster Information Base (CIB). As a consequence, certain cluster misconfigurations could remain undetected during routine operations.

With this enhancement, **pcs** has been updated to integrate Pacemaker's CIB validation tool into its workflow.

As a result, **pcs** now automatically performs a validation check and displays the results when you run the **pcs status**, **pcs cluster edit**, or **pcs cluster cib-push** commands.

Jira:RHEL-76060

pcs provides more detailed error messages for failed CIB updates

Previously, when a CIB update failed when using the **pcs cluster edit** or **pcs cluster cib-push** commands, the error message provided by Pacemaker was generic. It did not explain the specific reason for the failure, which made troubleshooting the invalid configuration difficult.

With this enhancement, **pcs** is updated to request a detailed validation check from Pacemaker upon a failed CIB push.

As a result, when a CIB update is rejected, **pcs** now displays a specific error message explaining what is wrong with the configuration.

Jira:RHEL-76059

A new pcs command is available for renaming a cluster

Previously, it was not possible to change the name of an existing cluster using **pcs** commands. Administrators had to perform a series of manual steps, which were complex and could lead to errors.

With this enhancement, the **pcs cluster rename** command has been introduced.

As a result, you can now easily change the name of an existing cluster. To rename your cluster, run the following command:

```
pcs cluster rename <new-name>
```

Jira:RHEL-76055

New fence agent for Nutanix AHV virtualization is now available

Previously, Red Hat High Availability Add-On did not provide a dedicated fence agent for Nutanix Acropolis Hypervisor (AHV) environments.

With this enhancement, the **fence_nutanix** agent is added.

As a result, you can now configure STONITH for cluster nodes running on the Nutanix AHV platform, enabling fully supported high-availability deployments.

Jira:RHEL-68321[1]

The pcs resource meta command is improved to support bundles and prevent guest node misconfiguration

Previously, the **pcs resource meta** command did not support managing meta attributes for bundle resources. Additionally, the command did not prevent users from incorrectly modifying the connection parameters of a guest node, which could lead to a misconfigured resource.

With this enhancement, the **pcs resource meta** command has been rewritten.

As a result, you can now use **pcs resource meta** to update meta attributes for bundle resources. In addition to this, when using the command on a guest node, it now prevents unintended changes to connection parameters, avoiding potential misconfigurations.

Jira:RHEL-35420

The IPaddr2 resource agent now detects network link failures

Before this update, the **IPaddr2** resource agent did not monitor the link state of the network interface. As a consequence, an **IPaddr2** resource continued to report success on a node even if the underlying

interface was in a **DOWN** or **LOWERLAYERDOWN** state, preventing the cluster from recovering the resource on another node.

With this release, the **IPaddr2** agent has been enhanced to check the interface's link status.

As a result, an **IPaddr2** resource correctly fails if its network interface goes down, allowing for a proper failover. You can disable this new default behavior by setting the **check_link_status=false** parameter in the resource configuration.

Jira:RHEL-7688[1]

The fence_aws agent supports immediate power-off

Previously, when the **fence_aws** agent performed an **off** or **reboot** action, it triggered a graceful shutdown of the instance. This introduced a delay in the fencing process, as the node was not powered off immediately.

With this update, a new **skip_os_shutdown** parameter has been added to the **fence_aws** agent. This parameter is enabled by default on Y-stream releases and disabled by default on Z-stream releases.

As a result, when **skip_os_shutdown** is set to **true**, the **fence_aws** agent bypasses the graceful shutdown and performs an immediate hard power-off of the instance.

Jira:RHEL-7601

4.12. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

The PostGIS extension is available for PostgreSQL 16

This enhancement adds the PostGIS extension to PostgreSQL 16. With this extension, PostgreSQL supports geographic objects, enabling spatial queries and analysis for Geographic Information System (GIS) applications, such as mapping, geolocation, and distance calculations within a relational database.

Jira:RHEL-81603[1]

4.13. COMPILERS AND DEVELOPMENT TOOLS

glibc now supports sched_setattr and sched_getattr for advanced scheduler options

Previously, **glibc** provided access to only a limited set of Linux scheduler options through functions defined in **<sched.h>**. This limitation required applications to use direct system calls or Linux kernel headers to access advanced scheduling features.

With this enhancement, the extensible scheduler configuration mechanism from **sched_setattr** and **sched_getattr** is now available through the **glibc <sched.h>** header file. This change includes support for additional scheduling policies, such as **SCHED_DEADLINE**.

As a result, applications can select from a wider range of scheduling options without relying on direct system calls or kernel-specific headers, improving portability and flexibility for developers.

Jira:RHEL-56627[1]

glibc pthread_gettid_np function added to libc_nonshared.a

Previously, there was no direct method to obtain the Linux task or thread ID (TID) from a glibc **pthread_t** handle. The newly implemented **pthread_gettid_np** function, declared in **<pthread.h>** when **_GNU_SOURCE** is defined, now allows applications that require TID, such as those using **sched_setattr**, to retrieve the TID value directly from a **pthread_t** handle.

As a result, applications can now use functions that expect a TID after obtaining it from a **pthread_t** handle, improving compatibility and simplifying thread management.

Jira:RHEL-83017

glibc fortification support added for inet_ntop and inet_pton

Previously, the **glibc** APIs **inet_ntop** and **inet_pton** did not include Source Fortification support, so the compiler was unable to catch some buffer errors before running the program.

With this update, attribute access annotations have been added to **inet_ntop** and **inet_pton**, enabling the compiler to warn about potential buffer misuse. The APIs are now covered by Source Fortification, improving their security and reliability.

Jira:RHEL-44920[1]

GDB now supports IBM's z17 CPU architecture

The **gdb** package is enhanced to support binaries that use new hardware instructions introduced with IBM's z17 CPU architecture. This update enables developers and system administrators to debug applications compiled for the latest IBM Z hardware on RHEL 9.7.

Jira:RHEL-50069[1]

GCC Toolset 15 is now available

With this update, **gcc-toolset-15** is now available in RHEL 9.7. The toolset includes the latest supported versions of GCC and related utilities, enabling developers to build, test, and deploy applications using up-to-date compiler technology.

Jira:RHEL-81741[1]

ELFv2 ABI support for -fpatchable-function-entry on ppc64le

Previously, the **-fpatchable-function-entry** option in **gcc** did not support the ELFv2 ABI on the ppc64le architecture, which caused NOP instructions to be generated in incorrect locations for that ABI. This issue prevented the correct use of the option when targeting ELFv2.

With this update, the **-fpatchable-function-entry** option can now be used on ppc64le to create programs for the ELFv2 ABI, ensuring NOPs are placed correctly and improving compatibility for users building on this platform.

Jira:RHEL-75806[1]

llvm-toolset rebased to LLVM 20

The **llvm-toolset** is updated to LLVM 20, delivering improved code generation, performance optimizations, and expanded language front-end and library support across C, C++, and Rust workflows. This rebase aligns dependent components in RHEL, including rebuilds for **rust**, **annobin**, **bcc**, **bpftrace**, **qt5-qttools**, and **mesa**. The build is validated with **llvm-20.1.8-3.el9**.

The notable changes are:

- Backend improvements, including fixes for the **ppc64le**.
- Optimizations and diagnostics enhancements in Clang and LLVM passes for general performance and reliability.
- Toolchain ecosystem refresh with coordinated package rebuilds for compatibility with LLVM 20.
- Continued deprecation of older targets consistent with upstream direction for ARM and MIPS in this stream.

Jira:RHEL-81006

Improved _r_debug extension support for debugging applications with multiple dynamic linker namespaces

The **glibc** package now includes the backported the **_r_debug** extension to support multiple namespaces. Previously, when attaching to running processes or analyzing core dumps, debuggers such as GDB could not display all loaded shared objects if the application used multiple namespaces with **dlmopen** or audit modules.

With this update, recent GDB versions can display shared objects across all dynamic linker namespaces, providing comprehensive debugging and analysis capabilities.

Jira:RHEL-101986[1]

Improved Exception Handling Performance in glibc

Before this update, exception handling in large applications was slow, impacting performance, particularly in environments with a high volume of users or frequent exceptions. This was due to the time spent in the **__dl_iterate_phdr** function, called from **_Unwind_Find_FDE**.

With this update, the exception handling algorithm in **glibc** has been improved to enhance exception processing speed. The update introduces new symbols to the ABI as part of GLIBC_2.35, including **__epoll_pwait2_time64**, **__memcmpeq**, **_dl_find_object**, **epoll_pwait2**, **posix_spawn_file_actions_addtcsetpgrp_np**, **posix_spawnattr_tcgetpgrp_np**, and **posix_spawnattr_tcsetpgrp_np**.

Jira:RHEL-93320

Hardening of glibc qsort behavior on memory allocation failure

When a memory allocation fails, the **qsort** and **qsort_r** functions of the **glibc** package use a heapsort fallback. This change improves handling of invalid comparison functions and makes performance more predictable if a memory allocation fails.

Because the fallback is not a stable sort, equal elements can appear in a different order. The C standard does not require stability.

Jira:RHEL-24168

gdb is rebased to version 16.3

This update of **gdb** to version 16.3 in RHEL 9.7 provides the following notable enhancements:

- Removed support for Intel MPX.

- Added support for tagged data pointers, including Intel's Linear Address Masking (LAM) and aarch64's Memory Tagging Extension (MTE).
- Enabled background DWARF reading for improved performance.
- Enhanced Intel Process Trace (**record btrace**):
  - Asynchronous event printing enabled with **set record btrace pt event-tracing**.
  - Ptwrite payloads can now be accessed in Python as **RecordAuxiliary** objects.
- Improved Python integration:
  - Stop events now include a **details** attribute, mirroring MI "*stopped" events.
  - **gdb.Progspace()** no longer creates objects directly; objects must be obtained with other APIs.
  - User-defined attributes can be added to **gdb.Inferior** and **gdb.InferiorThread** objects.
  - Introduced new event source: **gdb.tui_enabled**.
  - Added **gdb.record.clear**, which clears the current recording's trace data.
  - Added modules for handling missing objfiles and debug information.
  - New class **gdb.missing_debug.MissingDebugInfo** can be subclassed to handle missing debug information.
  - New attribute **gdb.Symbol.is_artificial**.
  - New constants for symbol lookup across multiple domains.
  - New function **gdb.notify_mi(NAME, DATA)** emits custom async notifications.
  - New attribute **gdb.Value.bytes** for reading and writing value contents.
  - Added **gdb.interrupt** to simulate a CTRL-C interrupt.
  - New attribute **gdb.InferiorThread.ptid_string** provides the target ID.
- Debug Adapter Protocol (DAP) changes:
  - Updated "scopes" request to include global variables and the last return value.
  - "launch" and "attach" requests can be used at any time, effective after "configurationDone".
  - "variables" request no longer returns artificial symbols.
  - Added "process" event and support for the "cancel" request.
  - "attach" request now supports specifying the program.
- Introduced new commands for styling, language frame mismatch warnings, missing objfile handlers, and function call timeouts.

- Enhanced and renamed several commands, including improved error handling for **disassemble** and renaming **set unwindonsignal** to **set unwind-on-signal**.
- Expanded remote packet support, including new packets for file status and memory fetch, and new stop reasons such as **clone**.
- Introduced per-thread event reporting options and address tagging checks.

Jira:RHEL-91381

AMD GPU pmda is now enabled for global GPU data collection

Before this update, the AMD GPU PMDA (a Performance Co-Pilot metrics agent) was not available in RHEL because the kernel lacked certain features required for full support.

With this update, users can now collect global GPU data on AMD GPUs in RHEL by using the **pcp-pmda-amdgpu** package.

Jira:RHEL-83154

Initial support for IBM Z z17 added to glibc

The dynamic loader in **glibc** is enhanced to support detecting IBM z17 CPUs or their specific features. As a result, any IBM z17-optimized libraries installed in the **/usr/lib64/glibc-hwcap/z17/** directory are loaded automatically on z17 systems. This update improves hardware compatibility and performance for IBM Z z17 platforms.

Jira:RHEL-50086[1]

Rust Toolset rebased to version 1.88.0

RHEL 9.7 is distributed with Rust Toolset in version 1.88.0. This update includes the following notable enhancements:

- Rust 2024 Edition is now stable. This is a major opt-in release that enables significant language changes and is the largest edition released to date.
- Leverage the 2024 Edition with **let** chains, allowing fluent **&&**-chaining of **let** statements within **if** and **while** conditions to reduce nesting and improve readability.
- For high-performance computing, when you enable target features, you can call multiple **std::arch** intrinsics directly in safe Rust, which gives you direct access to specific CPU features.
- **async** closures are now supported, providing first-class solutions for asynchronous programming. These closures allow borrowing from captures and properly express higher-ranked function signatures with the AsyncFn traits.
- Trait upcasting allows coercing a reference to a trait object to a reference of its supertrait, simplifying common patterns, especially with the **Any** trait.
- Cargo now automatically cleans its cache, removing old downloaded files not accessed in 1-3 months, which helps manage disk space.

Rust Toolset is a rolling Application Stream, and Red Hat only supports the latest version. For more information, see the Red Hat Enterprise Linux Application Streams Life Cycle document.

Jira:RHEL-81601

tzdata includes the NEWS file

With this update, the tzdata package includes its NEWS file with each release to provide precise descriptions of timezone data changes. As a result, you can review the changes in detail. Users can review the included NEWS file to understand what changed in the update.

Jira:RHEL-105043[1]

Metrics role now supports Apache Spark metric collection and export

Previously, users could not directly collect or export Apache Spark metrics using the metrics role. With this update, the **rhel-system-roles** package adds support to gather and update metrics from Apache Spark. Two new boolean parameters are introduced:

- **metrics_into_spark**: false This enables exporting metric values into Spark.
- **metrics_from_spark**: false This enables gathering metrics from Spark.

You can now both retrieve metrics from Spark and send metrics information into Spark, improving integration and monitoring capabilities for Spark workloads.

Jira:RHEL-78306

4.14. IDENTITY MANAGEMENT

ipa-healthcheck now warns about expiring certificates

With this update, the **ipa-healthcheck** tool now evaluates user-provided HTTP, DS, and PKINIT certificates for expiration and provides warnings 28 days prior to their expiration date. This is to prevent certificate expirations going potentially unnoticed, which can lead to downtime.

Jira:RHELDOCS-20303[1]

ansible-freeipa rebased to 1.15.1

The **ansible-freeipa** package, which provides modules and roles to manage Red Hat Identity Management (IdM) environments, has been rebased from version 1.13.2 to 1.15.1. The update includes the following enhancement:

- The **ansible-freeipa-collection** subpackage of **ansible-freeipa** is now compatible with the namespace and name of the **redhat.rhel_idm** collection provided by Red Hat Ansible Automation Hub (RH AAH). If you have installed the RPM collection subpackage, you can now run playbooks that reference the AAH roles and modules. Note that internally, the namespace and names from the RPM collection subpackage are used.

Jira:RHELDOCS-21029[1]

IdM now supports UIDs up to Linux maximum UID limit for legacy systems compatibility

With this update, you can now use User and Group IDs up to 4,294,967,293, or 2^32-1. This aligns IdM's maximum with the Linux UID limit and can be useful in rare cases where the standard IdM range, up to 2,147,483,647, is insufficient. Specifically, it enables IdM deployment alongside legacy systems that require the full 32-bit POSIX ID space.

WARNING

In standard deployments, IdM reserves the 2,147,483,648 - 4,294,836,223 range for subIDS. Using the 2^31 to 2^32-1 UID range requires disabling the subID feature and therefore conflicts with modern Linux capabilities.

To enable UIDs up to 2^32-1:

1. Disable the subordinate ID feature:

```
$ ipa config-mod --addattr ipaconfigstring=SubID:Disable
```

2. Remove any existing subordinate ID ranges:

```
$ ipa idrange-del <id_range>
```

3. On the IdM server, ensure the internal DNA plugin configuration is correctly removed:

```
# ipa-server-upgrade
```

4. Add a new local ID range that covers the 2^31 to 2^32-1 space. Ensure that you define RID bases for this new range so that IdM can generate SIDs properly for users and groups.

NOTE

You can only disable the subordinate ID feature if no subordinate IDs have been allocated yet.

Jira:RHEL-84277[1]

Healthcheck warns if krbLastSuccessfulAuth is enabled

Enabling the **krbLastSuccessfulAuth** setting in the **ipaConfigString** attribute can lead to performance issues if large numbers of users are authenticating at the same time. Therefore, it is disabled by default. With this update, **Healthcheck** displays a message if **krbLastSuccessfulAuth** is enabled, warning about the possible performance problems.

Jira:RHEL-4957

IdM-to-IdM migration now available

IdM-to-IdM migration, previously available as a Technology Preview, is now fully supported with this release. You can use the **ipa-migrate** command to migrate all IdM-specific data, such as SUDO rules, HBAC, DNA ranges, hosts, services, and more, from one IdM server to another. This can be useful, for example, when moving IdM from a development or staging environment into a production one.

Jira:RHELDOCS-19500[1]

samba rebased to version 4.22.4

The **samba** package has been updated to upstream version 4.22.4. This version provides bug fixes and enhancements, most notably the following:

- Samba supports Server message block version 3 (SMB3) directory leases. With this enhancement, clients can cache directory listings, which reduces network traffic and improves performance.
- Samba supports querying domain controller (DC) information by using TCP-based LDAP or LDAPS, as an alternative to the traditional UDP method on port 389. This enhancement improves compatibility with firewall-restricted environments. You can configure the protocol by using the **client netlogon ping protocol** parameter (default value: **CLADP**).
- The following configuration parameters are removed:
  - **nmbd_proxy_logon**: This setting was used to forward NetLogon authentication requests to a Windows NT4 primary domain controller (PDC) before Samba introduced its own NetBIOS over TCP/IP (NBT) server.
  - **cldap port**: Connectionless Lightweight Directory Access Protocol (CLDAP) always uses UDP port 389. Additionally, the Samba code did not use this parameter consistently, so the behavior was inconsistent.
  - **fruit:posix_rename**: This option of the **vfs_fruit** module is removed because it could result in problems with Windows clients. As a possible workaround to prevent the creation of **.DS_Store** files on network mounts, use the **defaults write com.apple.desktopservices DSDontWriteNetworkStores true** command on MacOS.

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.

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

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

Jira:RHEL-89873

389-ds-base rebased to version 2.7.0

The **389-ds-base** package has been updated to version 2.7.0.

Jira:RHEL-80163

dsctl healthcheck now warns about creating a substring index on the membership attribute

An entry that contains a membership attribute is usually a group with many members. When changing the value set, substring index is very expensive even for a minor change like deleting a single member. Now, when you add the substring index type, **dsctl healthcheck** warns about possible high cost of substring index on membership attributes and displays the following error message:

DSMOLE0002. If the substring index is configured for a membership attribute, the removal of a member from the large group can be slow.

Jira:RHEL-81141[1]

Custom matching rules in the Attribute Uniqueness plug-in to search uniqueness attributes

With this update, in Attribute Uniqueness plug-in configuration, you can specify a matching rule for the attribute you want to enforce uniqueness on. For example, when you want to override the attribute's syntax from **case exact** or **case ignore**.

Specify attributes and their matching rules in the plugin configuration, as follows:

```
uniqueness-attribute-name: <attribute>:<Matching rule OID>:
```

Before this update, if you used the attribute **cn** with a **case exact** syntax, the Attribute Uniqueness plug-in could not find a matching value if the case was different between the two values being compared. Now you can set the matching rule and make it **case ignore** and the plug-in will see that the values match:

```
uniqueness-attribute-name: cn:caseIgnoreMatch:
```

Jira:RHEL-109034[1]

cockpit-session-recording rebased to 20-1.el9

The **cockpit-session-recording** package, which records user sessions that are conducted through the Cockpit web interface, is rebased to upstream version 20-1.el9. The package has been migrated to PatternFly 6 user interface system design.

Jira:RHEL-96905

ACME server adds support for the ES256 signature algorithm

Previously, the Automatic Certificate Management Environment (ACME) server did not support the ES256 signature algorithm for JSON Web Key (JWK) validation. This lack of support prevented certain clients, such as the Caddy web server, from successfully obtaining certificates.

With this update, the ACME server has been enhanced to support the ES256 signature algorithm for JWK validation.

As a result, the server can interoperate with clients that use ES256, such as the Caddy web server, allowing them to successfully obtain certificates and establish secure HTTPS communication.

Jira:RHEL-98719

HSM is now fully supported in IdM

Hardware Security Modules (HSM) are now fully supported in Identity Management (IdM). You can store your key pairs and certificates for your IdM Cerificate Authority (CA) and Key Recovery Authority (KRA) on an HSM. This adds physical security to the private key material.

IdM relies on the networking features of the HSM to share the keys between machines to create replicas. The HSM provides additional security without visibly affecting most IdM operations. When using low-level tooling the certificates and keys are handled differently but this is seamless for most users.

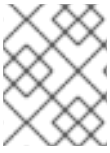

NOTE

Migration of an existing CA or KRA to an HSM-based setup is not supported. You need to reinstall the CA or KRA with keys on the HSM.

You need the following:

- A supported HSM.
- The HSM Public-Key Cryptography Standard (PKCS) #11 library.
- An available slot, token, and the token password.

To install a CA or KRA with keys stored on an HSM, you must specify the token name and the path to the PKCS #11 library. For example:

```
ipa-server-install -r EXAMPLE.TEST -U --setup-dns --allow-zone-overlap --no-forwarders -N --auto-reverse --random-serial-numbers -—token-name=HSM-TOKEN --token-library-path=/opt/nfast/toolkits/pkcs11/libcknfast.so --setup-kra
```

Jira:RHELDOCS-21376[1]

4.15. DESKTOP

OpenGL and Vulkan are supported by default in Toolbox containers based on UBI

OpenGL and Vulkan now work by default inside Toolbox containers created from updated UBI-based toolbox images, matching the behavior on RHEL Workstation hosts. This includes only the free software drivers provided by Mesa, not proprietary ones like NVIDIA.

Toolbx containers aim to replicate the RHEL Workstation environment. Previously, users had to manually install Mesa-related packages to enable OpenGL and Vulkan support, which was not intuitive or documented.

As a result, OpenGL and Vulkan applications can run inside Toolbox containers without additional configuration, improving usability and consistency with the host system.

Jira:RHEL-84787

Low Disk Space notifications include a mount point in the web console

The **Low Disk Space** notifications include the mount point when multiple volumes have the same name. This enhancement reduces ambiguity about which specific file system requires more space.

Jira:RHEL-11910[1]

4.16. THE WEB CONSOLE

cockpit rebased to version 344

The **cockpit** packages have been rebased to version 344, which provides many improvements and fixes compared to version 334 in RHEL 9.6, most notably:

- Improved UI to the new style based on the PatternFly 6 design system.
- Added support for the SMART (Self-Monitoring, Analysis and Reporting Technology) standard and the Stratis 3.8+ pool format in the Storage component.
- Improved graphical VNC, control VNC, and serial consoles in the Virtual machines component.
- Added support for IPv6 addresses for WireGuard VPNs in the Networking component.

- All web console pages can be branded through the **branding.css** style-sheet file.

Jira:RHEL-87397

new subpackage: cockpit-ws-selinux

The SELinux policy for the **cockpit_ws** processes is provided in a separate subpackage **cockpit-ws-selinux**. This prevents the RHEL web console from failing when run on a system without SELinux installed, because the package manager installs the **selinux_policy** packages as dependencies. See the **cockpit_ws_selinux(8)** man page on your system for more information.

Jira:RHEL-92062

4.17. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The ad_integration RHEL system role can control the SSSD domain section naming and consolidate duplicates

With this update, users can control the name of the section used in the SSSD config file for the domain or realm-specific settings, as managed by the **ad_dyndns_update** and **ad_integration_sssd_custom_settings** parameters. By default, the **ad_integration** role uses the lower case of the **ad_integration_realm** variable. However if users want to use the actual case of **ad_integration_realm**, users can use a new option **ad_integration_sssd_realm_preserve_case = true** to preserve the case of the realm. This may leave the SSSD config file with multiple sections for the realm. Use the new **ad_integration_sssd_remove_duplicate_sections** setting to consolidate all of the settings from the multiple sections into the chosen section. As a result, the **ad_integration** system role can manage domain and realm sections in the SSSD config file correctly.

Jira:RHEL-99089[1]

The journald RHEL system role can monitor disk space

With this update, you can configure the **SystemKeepFree** option in the **journald.conf** journal service to set a maximum size for the system journal. This improves overall system stability and performance. As a result, you can use the **journald_system_keep_free** variable to configure size limit. The value is specified in megabytes. There is no default value - by default, it will use the **journald** default value.

Jira:RHEL-95874[1]

metrics role supports enabling additional PCP domains

With this update, the **rhel-system-roles** package introduces the **metrics_optional_domains** variable in the **metrics** RHEL system role. Users can specify a list of additional PCP domains to be activated, in addition to those that are automatically managed by the **metrics** role. As a result, users can enable the domains they require for their specific use cases, improving flexibility in data collection and monitoring.

Jira:RHEL-104659[1]

Introduced a variable MaxRetention to configure the maximum retention parameter for journald

With this update, users can configure the maximum retention parameter for **journald**, enabling time-based deletion of journal files. This enhancement provides flexibility in managing log data according to specific data retention policies, allowing both time-based log deletion and size-based deletion. It helps

with compliance with data retention requirements and improves overall system performance by preventing excessive log storage.

Jira:RHEL-102637[1]

The podman role generates all TOML compliant configuration file

Before this update, the current Jinja-based formatter did not support many TOML features, including tables and inline tables, which were required to configure all aspects of **podman**. With this enhancement, all features of TOML are supported by using a true TOML formatter instead of a simple Jinja template. As a result, the **podman** role can generate any TOML compliant configuration file that **podman** can use.

The **podman** role needs to preserve certain features of the old formatter. Therefore, the TOML formatter is disabled by default. For the particular use cases that you need to use the old formatter for and information about how you can convert your inventory data in order to use the new and improved formatter, see the README file.

To use the new TOML formatter in all cases, set the **podman_use_new_toml_formatter** to **true**:

```
podman_use_new_toml_formatter: true
```

Jira:RHEL-84930

The firewall RHEL system role now supports including other services

With this enhancement, you can include other services when you use the **firewall** RHEL system role to create **firewalld** service definitions. For example, you can create a service **webserver** that includes the **http** and **https** services. If you then enable the **webserver** service, **firewalld** open the ports defined in **http** and **https** services. For further details, see Creating a custom firewalld service by using the firewall RHEL system role.

Jira:RHEL-84951

Ability to configure the default kernel in rhel-system-roles

Previously, users could not specify which kernel should be set as the default during system boot. This limitation prevented administrators from easily managing the default kernel selection during automation.

With this update, the **rhel-system-roles** package allows configuring the default bootloader kernel using a new **default** option. Users can now designate a single kernel as the default by setting the **default** boolean parameter in kernel settings. The system validates that only one kernel can be marked as default, and applies the selection using **grubby --set-default** as required.

This enhancement improves flexibility and simplifies automation when managing kernel versions in RHEL.

Jira:RHEL-87579

Metrics role now supports Apache Spark metric collection and export

Previously, users could not directly collect or export Apache Spark metrics using the metrics role. With this update, the **rhel-system-roles** package adds support to gather and export metrics from Apache Spark. Two new boolean parameters are introduced:

- **metrics_into_spark**: false This enables exporting metric values into Spark.

- **metrics_from_spark**: false This enables gathering metrics from Spark.

You can now both retrieve metrics from Spark and send metrics information into Spark, improving integration and monitoring capabilities for Spark workloads.

Jira:RHEL-17564

Enables IPv4-only operation for the chronyd service when using the rhel-system-roles.timesync role

With this update, users can customize the **chronyd** configuration when IPv6 is disabled on a node. The enhancement provides two options: add a setting to the **timesync** role to disable IPv6, or pass a parameter to set the OPTIONS value for **chronyd**. These options enable IPv4-only operation for the **chronyd** service when using the **rhel-system-roles.timesync** role. This improves time synchronization accuracy and stability for environments where IPv6 is disabled.

Jira:RHEL-85079

4.18. VIRTUALIZATION

virtio-mem is available on IBM Z

With this update, **virtio-mem**, a paravirtualized memory device, can be used on IBM Z hardware. By using **virtio-mem**, you can dynamically add or remove host memory in virtual machines.

Jira:RHEL-72976[1]

New command for IBM Z hosts: virsh hypervisor-cpu-models

This update introduces the **virsh hypervisor-cpu-models** command. You can use this command on the IBM Z architecture to display which CPU models your hypervisor recognizes.

Jira:RHEL-11435[1]

Performance-enhanced PCI translation for IBM Z guests

With this update, virtual machines (VMs) on IBM Z hosts can use identity-mapped direct memory access (DMA) for PCI devices. This feature significantly improves the performance of PCI device passthrough. Note that to use the feature, your system must be configured as follows:

- The **iommu.passthrough=1** parameter must be set up on the kernel command line of the VM.
- The VM must have fully NUMA-pinned memory.
- The RHEL host system must not be using logical partitioning (LPAR).

Jira:RHEL-11431[1]

New features for virtual machines on 64-bit ARM hosts

The following features are now supported for virtual machines on RHEL hosts that use the 64-bit ARM architecture(**aarch64**):

- Live snapshots
- Pre-copy migration with the following options:

    - TLS encryption and XBZRLE compression
    - Dirty rate monitoring
    - Auto-converge
- Multi-FD migration with the following options:
    - TLS encryption and XBZRLE compression
    - Auto-converge
    - Zero-copy
- Post-copy migration with the following options:
    - TLS encryption and XBZRLE compression
    - Recovery
    - Preemption
- Live migration with **virtiofs**
- Backward migration from RHEL 10.1 to RHEL 9.7

Jira:RHELDOCS-20781[1]

4.19. RHEL IN CLOUD ENVIRONMENTS

OTel collector on RHEL supports TPM device

The OpenTelemetry (OTel) Collector on RHEL supports the Trusted Platform Module (TPM) device. With this feature, OTel Collector can read transport layer security (TLS) certificates from the TPM device.

Jira:RHELDOCS-20446[1]

Enhanced automatic registration for eligible RHEL images

With this update, RHEL instances based on eligible images from eligible marketplaces automatically receive content and updates from Red Hat content delivery network (CDN) instead of the Red Hat Update Infrastructure (RHUI). The RHUI repositories are turned off by default.

This ensures automatic access to latest updates for users of subscribed RHEL instances.

For additional details, see Understanding auto-registration.

Jira:RHELDOCS-21241[1]

New package: azure-vm-utils

This update adds the **azure-vm-utils** package, which provides a collection of utilities and **udev** rules to optimize the experience of using RHEL 9 as a guest operating system on Microsoft Azure.

Jira:RHEL-88789[1]

RHEL is available on Azure confidential VMs

You can create and run RHEL confidential virtual machines (CVMs) on Microsoft Azure by using RHEL CVM images. The images support full disk encryption through the Confidential OS disk encryption feature in Azure.

Jira:RHELPLAN-139800[1]

Enhanced automatic registration for eligible RHEL images

When purchasing certain eligible cloud marketplace subscriptions for RHEL 9.6 or later and for RHEL 10.0 or later, an improved version of the auto-registration function is available.

With the enhanced auto-registration, any RHEL instances on the eligible marketplaces will be automatically registered to Red Hat and automatically receive content updates from Red Hat Update Infrastructure (RHUI) after you establish a trusted connection between your Red Hat account and your account for the respective cloud platform, even if you did not have the trusted connection when you set launched the instance.

For additional details, see Understanding auto-registration.

Jira:RHELDOCS-19664[1]

4.20. SUPPORTABILITY

sos now collects the Satellite metrics file for improved support diagnostics

The **foreman-installer** plugin of **sos** now collects the **satellite_metrics.yml** file located at **/var/lib/foreman-maintain/** directory. It provides insight into which features of Satellite are in use and in what scale.

Jira:RHEL-71825

4.21. CONTAINERS

A new rhel9/valkey-8 container image is generally available in RHEL

The newly available **rhel9/valkey-8** container image allows atomic operations and supports various data types like strings, hashes, lists, sets, and sorted sets. The image offers high performance because of its in-memory dataset, which can be persisted to disk or by appending commands to a log.

Jira:RHELDOCS-20639[1]

Improved support for reproducible container builds

Reproducible builds ensure that a given set of inputs consistently generates the same output. This enhancement addresses several factors that previously complicated reproducibility in container image builds. While using **-source-date-epoch** and **-rewrite-timestamp** improves the reproducibility of builds and better aligns with common practices like setting and looking for **$SOURCE_DATE_EPOCH**, it cannot guarantee complete reproducibility.

Jira:RHEL-88521

New artifact endpoints for Podman RESTFUL API

Podman RESTFUL API includes new artifact endpoints, enabling programmatic management of OCI artifacts. This enhancement simplifies integration of OCI artifact operations into existing systems and scripts.

Jira:RHEL-88472

The Container Tools packages have been updated

The updated Container Tools RPM meta-package, which contains the Podman, Buildah, Skopeo, **crun**, and **runc** tools, is available. The Buildah package has been updated to version v1.41.0, and Skopeo has been updated to version 1.20.0.

Podman release v5.6 contains the following notable bug fixes and enhancements over the previous version:

- A new set of commands for managing Quadlets has been added as **podman quadlet install** (install a new Quadlet for the current user), **podman quadlet list** (list installed Quadlets), **podman quadlet print** (print the contents of a Quadlet file), and **podman quadlet rm** (remove a Quadlet).
- The **podman kube play** command can restrict container execution to specific CPU cores and specific memory nodes using the **io.podman.annotations.cpuset/$ctrname** and **io.podman.annotations.memory-nodes/$ctrname** annotations.
- The **podman kube play** command supports the **lifecycle.stopSignal** field in Pod YAML, allowing the signal used to stop containers to be specified.
- The **podman volume import** and **podman volume export** commands are available in the remote Podman client.
- The **podman volume create** command accepts two new options, **--uid** and **--gid**, to set the UID and GID the volume will be created with.
- The **podman secret create** command has a new option, **--ignore**, causing the command to succeed even if a secret with the given name already exists.
- The **podman pull** command has a new option, **--policy**, to configure pull policy.
- The **podman update** command has a new option, **--latest**, to update the latest container instead of specifying a specific container.
- A full set of API endpoints for interacting with artifacts has been added, including inspecting artifacts (**GET /libpod/artifacts/{name}/json**), listing all artifacts (**GET /libpod/artifacts/json**), pulling an artifact (**POST /libpod/artifacts/pull**), removing an artifact (**DELETE /libpod/artifacts/{name}**), adding an artifact (or appending to an existing artifact) from a tar file in the request body (**POST /libpod/artifacts/add**), pushing an artifact to a registry (**/libpod/artifacts/{name}/push**), and retrieving the contents of an artifact (**GET /libpod/artifacts/{name}/extract**).
- A new command has been added, **podman artifact extract**, to copy some or all of the contents of an OCI artifact to a location on disk.
- The **--mount** option to **podman create**, **podman run**, and **podman pod create** supports a new mount type, **--mount type=artifact**, to mount OCI artifacts into containers.

- The **podman artifact add** command features two new options, **--append** to add new files to an existing artifact, and **--file-type** to specify the MIME type of the file added to the artifact.
- The **podman artifact rm** command features a new option, **--all**, to remove all artifacts in the local store.
- The **podman kube generate** and **podman kube play** commands supports a new annotation, **io.podman.annotation.pids-limit/$containername**, preserving the PID limit for containers across **kube generate** and **kube play**.
- Quadlet **.container** units support three new keys, **Memory=** (set maximum memory for the created container), **ReloadCmd** (execute a command via systemd **ExecReload**), and **ReloadSignal** (kill the container with the given signal via systemd **ExecReload**).
- Quadlet **.container**, **.image**, and **.build** units support two new keys, **Retry** (number of times to retry pulling image on failure) and **RetryDelay** (delay between retries).
- Quadlet **.pod** units support a new key, **HostName=**, to set the pod's hostname.
- Quadlet files support a new option, **UpheldBy**, in the **Install** section, corresponding to the systemd **Upholds** option.
- The names of Quadlet units specified as systemd dependencies are automatically translated, for example **Wants=my.container** is valid.

For more information about notable changes, see upstream release notes.

Jira:RHEL-88464

The ADD and COPY instructions now support the --link option

Buildah and Podman now support the **--link** flag for ADD and COPY instructions in Containerfiles, which causes the new content to be added as its own layer in the built image.

Jira:RHEL-88307

New container images are available

The new container images are listed in the Red Hat Ecosystem Catalog:

- **ubi-stig**: the Universal Base Image with STIG hardening as a secure foundation for containerized applications, middleware, and utilities.
- **valkey-8**: an advanced key-value store available as a container, uses an in-memory dataset to achieve its outstanding performance. It is often referred to as a data structure server because keys can contain strings, hashes, lists, sets, and sorted sets.
- **gcc-toolset-15-toolchain**: a base image with essential libraries and tools used to build C and C++ applications.
- **nodejs-24**: provides a base platform for building and running various Node.js 24 applications and frameworks. It is built on Chrome's JavaScript runtime, it facilitates fast, scalable network applications through an event-driven, non-blocking I/O model, ideal for data-intensive real-time distributed applications.
- **nodejs-24-minimal**: provides a base platform for running various Node.js 24 applications and frameworks. It is built on Chrome's JavaScript runtime, it facilitates fast, scalable network

applications through an event-driven, non-blocking I/O model, ideal for data-intensive real-time distributed applications.

- **dotnet-100**, **dotnet-100-aspnet**, **dotnet-100-runtime**: The .NET 100 images, including base, ASP.NET, and runtime versions, are now available.

Jira:RHELDOCS-21211[1]

RHEL image mode supports creating root-level directories and symlinks at runtime

With this release, you can use RHEL image mode to create root-level directories and symbolic links after system deployment, then return the filesystem to read-only mode. As a result, you can use a single base image across multiple deployment environments with different file system requirements.

Jira:RHELDOCS-21230[1]

bootc-image-builder uses the local container storage by default

With this release, the **bootc-image-builder** tool operates in local mode by default, which means it no longer pulls container images from remote registries. To build disk images, you must pre-load the base bootc container image in the local container registry of the system before building disk images. If you have existing workflows that relied on automatic image pulling, you must update them. This change improves security by reducing external network dependencies during the build process.

Jira:RHELDOCS-21218[1]

4.22. RHEL LIGHTSPEED

The command-line assistant supports image mode for RHEL

With this enhancement, you can customize your Containerfile to include the **command-line-assistant** package, create a disk image from a container image, and boot a system with that image. As a result, the system image has the command-line assistant preinstalled, and you can use it after you register your system with **subscription-manager**.

Jira:RHELDOCS-20546[1]

The command-line assistant context limit increased to 32KB input

Before this update, the command-line assistant had a 2KB input context limit, causing it to fail when input exceeded this limit. As a consequence, user experience was limited, preventing thorough log analysis due to the 2KB input context limit. With this release, the command-line assistant input context limit has been increased from 2KB to 32KB. As a result, the command-line assistant now supports larger input contexts, enabling better log analysis and potential issue detection.

Jira:RHELDOCS-20421[1]

The command-line assistant for RHEL Lightspeed has better error handling and exit codes

With this enhancement, the command-line assistant brings better error handling and exit codes, such as:

- Output different error messages based on different types of errors that can occur during CLA runtime.
- Try to output an error message that corresponds to the actual cause of the error, and log it.
- Implement different exit codes based on different types of issues.

Jira:RHELDOCS-21313[1]

Command-line assistant -w option displays current output

Before this update, when you tried to use the **-w** option without the current enable-capture mode, the command-line assistant incorrectly displayed output from an earlier session. With this update, the terminal capture log file is actively verified before outputting from the **-w** option. As a result, the mentioned problem is fixed, and the displayed output is accurate.

Jira:RHELDOCS-21315[1]

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 9.7. 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

arm64.nompam=

[ARM64] Unconditionally disable Memory Partitioning And Monitoring support.

indirect_target_selection=

[X86,Intel] Mitigation control for Indirect Target Selection (**ITS**) bug in Intel CPUs. Updated microcode is also required for a fix in **IBPB**.

Possible values:

on – Enable mitigation (default). **off** – Disable mitigation. **force** – Force the **ITS** bug and deploy default mitigation. **vmexit** – Only deploy mitigation if the CPU is affected by guest/host isolation part of **ITS**. **stuff** – Deploy **RSB**-fill mitigation when **retpoline** is also deployed. Otherwise, deploy the default mitigation.

See Documentation /**admin-guide/hw-vuln/indirect-target-selection.rst**

pcie.notph

[PCIE] If the **PCIE_TPH** kernel configuration parameter is enabled, this kernel boot option can be used to disable PCIe TLP Processing Hints support system-wide.

rcutree.csd_lock_suppress_rcu_stall=

[KNL] Do only a one-line RCU CPU stall warning when there is an ongoing too-long **CSD**-lock wait.

rcuscale.kfree_by_call_rcu=

[KNL] In kernels built with **CONFIG_RCU_LAZY=y**, test **call_rcu()** instead of **kfree_rcu()**.

rcuscale.kfree_mult=

[KNL] Instead of allocating an object of size **kfree_obj**, allocate one of **kfree_mult * sizeof(kfree_obj)**. Defaults to **1**.

rcuscale.scale_type=

[KNL] Specify the RCU implementation to test.

rcutorture.stall_cpu_repeat=

[KNL] Number of times to repeat the stall sequence, so that **rcutorture.stall_cpu_repeat=3** will result in four stall sequences.

refscale.lookup_instances=

[KNL] Number of data elements to use for the forms of **SLAB_TYPESAFE_BY_RCU** testing. A negative number is negated and multiplied by **nr_cpu_ids**, while zero specifies **nr_cpu_ids**.

smp.panic_on_ipistall=

[KNL] If a **csd_lock_timeout** extends for more than the specified number of milliseconds, panic the system. By default, let **CSD**-lock acquisition take as long as they take. Specifying **300000** for this value provides a 5-minute timeout.

spectre_bhi=

[X86] Control mitigation of Branch History Injection (**BHI**).

on – (default) Enable the HW or SW mitigation as needed. This protects the kernel from both syscalls and VMs. **vmexit** – On systems which don't have the HW mitigation available, enable the SW mitigation on **vmexit** ONLY. On such systems, the host kernel is protected from VM-originated **BHI** attacks, but may still be vulnerable to syscall attacks. **off** – Disable the mitigation.

tsa=

[X86] Control mitigation for Transient Scheduler Attacks on AMD CPUs. Search the following in your favourite search engine for more details:

Technical guidance for mitigating transient scheduler attacks. **off** – disable the mitigation **on** – enable the mitigation (default) **user** – mitigate only user/kernel transitions **vm** – mitigate only guest/host transitions

Removed kernel parameters

clocksource.max_cswd_read_retries=

[KNL] Number of **clocksource_watchdog()** retries due to external delays before the clock will be marked unstable. Defaults to two retries, that is, three attempts to read the clock under test.

disable_cpu_apicid=

[X86,APIC,SMP] Format: **<int>** The number of initial **APIC ID** for the corresponding CPU to be disabled at boot, mostly used for the kdump 2nd kernel to disable **BSP** to wake up multiple CPUs without causing system reset or hang due to sending **INIT** from **AP** to **BSP**.

Changed kernel parameters

nohz_full=

[KNL] Disable the tick when a single task runs as well as disabling other kernel noises like having **RCU** callbacks offloaded. This is equivalent to the **nohz_full** parameter. A residual **1Hz** tick is offloaded to workqueues, which you need to affine to housekeeping through the global **sysfs** interface.

mce=

[X86-64] See Documentation /**arch/x86/x86_64/boot-options.rst**.

mem_encrypt=

[X86-64] See Documentation /**virt/kvm/x86/amd-memory-encryption.rst** for details on when memory encryption can be activated.

mitigations=

[ALL] Selecting **mitigations=off** is equivalent to also turning off the following:

If **nokaslr** then **kpti=0 [ARM64]gather_data_sampling=off [X86]indirect_target_selection=off [X86]kvm.nx_huge_pages=off [X86]l1tf=off [X86]mds=off [X86]meltdown=off [X86]mmio_stale_data=off [X86]pcid=off [X86]pti=off [X86]spectre_v1=off [X86]spectre_v2=off [X86]tsx=off [X86]tsx_async_abort=off [X86]uhi=off [X86]**

pci=config_acs=

[PCI] Format: **<ACS flags>@<pci_dev>[; ...]**

Each bit value:

0 – force disabled **1** – force enabled **x** – unchanged For example, **pci=config_acs=10x@pci:0:0** would configure all devices that support **ACS** to enable **P2P Request Redirect**, disable **Translation Blocking**, and leave **Source Validation** unchanged from whatever power-up or firmware set it to.

NOTE

This may remove isolation between devices and may put more devices in an IOMMU group.

pirq=

[SMP,APIC] See Documentation **/arch/x86/i386/IO-APIC.rst**.

prot_virt=

[S390] Enable hosting protected virtual machines isolated from the hypervisor (if hardware supports that). If enabled, the default kernel base address might be overridden even when Kernel Address Space Layout Randomization is disabled. Format: **<bool>**

sev=

[X86-64] See Documentation **/arch/x86/x86_64/boot-options.rst**.

spectre_v2_user=

[X86] Control mitigation of Spectre variant 2 for user space. Selecting **on** will also enable the mitigation against user space to user space task attacks. Selecting specific mitigation does not force enable user mitigations. Selecting **off** will disable both the kernel and the user space protections.

rcutorture.stall_cpu_irqsoff=

[KNL] Disable interrupts while stalling if set, but only on the first stall in the set.

New sysctl parameters

timer_migration

When set to a non-zero value, attempt to migrate timers away from idle CPUs to allow them to remain in low power states longer. Default: **1**

CHAPTER 6. DEVICE DRIVERS

6.1. NEW DRIVERS

Table 6.1. Character device drivers

Description	Name	Limited to architectures
SNP SVSM vTPM Driver	tpm_svsm	AMD and Intel 64-bit architectures

Table 6.2. CPU frequency drivers

Description	Name	Limited to architectures
Virtual cpufreq driver	virtual-cpufreq	64-bit ARM architecture

Table 6.3. DMA drivers

Description	Name	Limited to architectures
AMD AE4DMA driver	ae4dma	AMD and Intel 64-bit architectures
AMD PassThru DMA driver	ptdma	AMD and Intel 64-bit architectures

Table 6.4. Firmware control drivers

Description	Name	Limited to architectures
fwctl device firmware access framework	fwctl	
mlx5 ConnectX fwctl driver	mlx5_fwctl	

Table 6.5. Graphics drivers and miscellaneous drivers

Description	Name	Limited to architectures
Chrontel ch7006 TV encoder driver	ch7006	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Cirrus driver for QEMU emulated device	cirrus-qemu	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
In-kernel DRM clients	drm_client_lib	
NXP Semiconductors TDA998X HDMI Encoder	tda998x	64-bit ARM architecture
Quirks for panel backlight overrides	drm_panel_backlight_quirks	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Silicon Image sil164 TMDS transmitter driver	sil164	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures

Table 6.6. HID drivers

Description	Name	Limited to architectures
Intel® Intel THC Hardware Driver	intel-thc	AMD and Intel 64-bit architectures
Intel® QuickI2C Driver	intel-quicki2c	AMD and Intel 64-bit architectures
Intel® QuickSPI Driver	intel-quickspi	AMD and Intel 64-bit architectures

Table 6.7. Media drivers

Description	Name	Limited to architectures
Conexant cx231xx based USB video device driver - 0.0.3	cx231xx	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Conexant CX25840 audio/video decoder driver	cx25840	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
cx23415/6/8 driver	cx2341x	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
device driver for various TV and TV+FM radio tuners	tuner	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
i2c Hauppauge eeprom decoder driver	tveeprom	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures

Table 6.8. Network drivers

Description	Name	Limited to architectures
Intel® MLD wireless driver for Linux	iwlmld	64-bit ARM architecture, AMD and Intel 64-bit architectures
Socket CAN device driver for Geschwister Schneider Technologie- Entwicklungs- und Vertriebs UG. USB2.0 to CAN interfaces and bytewerk.org candleLight USB CAN interfaces.	gs_usb	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures

Table 6.9. Platform drivers

Description	Name	Limited to architectures
AMD 3D V-Cache Performance Optimizer Driver	amd_3d_vcache	AMD and Intel 64-bit architectures
Intel Extended Capabilities auxiliary bus driver	intel-vsec	AMD and Intel 64-bit architectures
Intel Oaktrail Platform ACPI Extras - 0.4ac1	intel-oaktrail	AMD and Intel 64-bit architectures
Intel On Demand (SDSi) driver	intel-sdsi	AMD and Intel 64-bit architectures
Intel TPMI enumeration module	intel-vsec_tpmi	AMD and Intel 64-bit architectures
Intel TPMI PLR Driver	intel-plr_tpmi	AMD and Intel 64-bit architectures
ISH ISHTP eclite client opregion driver	intel-ishtp_eclite	AMD and Intel 64-bit architectures
lis3lv02d i2c-client instantiation for ACPI SMO88xx devices	dell-lis3lv02d	AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
TPMI Power Domains Mapping	intel-tpmi_power_domains	AMD and Intel 64-bit architectures

Table 6.10. Power management domain drivers

Description	Name	Limited to architectures
ARM SCMI power domain driver	scmi_pm_domain	64-bit ARM architecture
ARM SCPI power domain driver	scpi_pm_domain	64-bit ARM architecture

Table 6.11. USB drivers

Description	Name	Limited to architectures
Thunderbolt3 USB Type-C Alternate Mode	typec_thunderbolt	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures

Table 6.12. Virtio drivers

Description	Name	Limited to architectures
Virtio-mem driver	virtio_mem	IBM Z

6.2. UPDATED DRIVERS

Table 6.13. Accelerator driver updates

Description	Name	Current version	Limited to architectures
Driver for Intel NPU (Neural Processing Unit)	intel_vpu	1.0.0 (5.14.0-611.5.1.el9_7.x86_64)	AMD and Intel 64-bit architectures

Table 6.14. Platform driver updates

Description	Name	Current version	Limited to architectures
Driver for updating BIOS image on DELL systems	dell_rbu	3.3	AMD and Intel 64-bit architectures

Table 6.15. Storage driver updates

Description	Name	Current version	Limited to architectures
Broadcom MegaRAID SAS Driver	megaraid_sas	07.734.00.00-rc1	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Cisco FCoE HBA Driver	fnic	1.8.0.2	AMD and Intel 64-bit architectures
Driver for Microchip Smart Family Controller	smartpqi	2.1.34-035	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Emulex LightPulse Fibre Channel SCSI driver	lpfc	0:14.4.0.9	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
LSI MPT Fusion SAS 3.0 Device Driver	mpt3sas	52.100.00.00	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
MPI3 Storage Controller Device Driver	mpi3mr	8.15.0.5.50	

CHAPTER 7. BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 9.7 that have a significant impact on users.

7.1. INSTALLER AND IMAGE CREATION

Installation no longer fails if a VDO logical volume is present

Before this update, installing RHEL failed when users attempted to remove a pre-existing Logical Volume Manager Virtual Data Optimizer (LVM VDO) volume on systems without the **dm_vdo** kernel module. With this update, installation succeeds when removing an LVM VDO volume on systems without VDO support.

Jira:RHEL-8008[1]

Installer now respects the BOOTIF boot argument

Previously, the RHEL installer ignored the **BOOTIF=<MAC>** boot argument and activated all the available network interfaces. With this fix, the installation program now properly processes the **BOOTIF** argument and ensures that only the designated network device is activated during the installation process.

Jira:RHEL-78272[1]

7.2. SECURITY

SSH connection fail no longer displays verbose help message

Before this update, when SSH connection failed, a message with common SSH errors and a link to Red Hat help was displayed. As a consequence, the help message in the error output broke user scripts and automation. With this update, the help message displays only when SSH is run with log level **debug1** or higher. As a result, the error output does not include any unexpected messages by default.

Jira:RHEL-104580[1]

OpenSC avoids memory freeing before dereferencing

Before this update, dereferencing would free members when OpenSC was reading public keys. This caused unpredictable behavior of the values stored in the memory. This update avoids freeing the memory before dereferencing. As a result, OpenSC correctly handles reading public keys.

Jira:RHEL-96029

fapolicyd no longer causes the RPM database to crash with repeated updates

Before this update, repeated updates of the RPM database when **fapolicyd** was in enforcing mode caused a bus error (SIGBUS), which caused the RPM database to terminate unexpectedly. With this release, fapolicyd SIGBUS protection for RPM database updates has been improved. As a result, the RPM database no longer crashes when repeatedly updating it with **fapolicyd** enabled.

Jira:RHEL-63090[1]

fapolicyd-cli --file add no longer fails when processing non-regular files

Before this update, the **fapolicyd-cli --file add** command failed to add directories containing non-regular files, such as sockets, to the trust database. With this update, the problem is resolved and **fapolicyd-cli --file add** no longer fails in the described scenario.

Jira:RHEL-69136

fapolicyd no longer fails to identify user accounts from a network source

Before this update, due to an incorrect security policy configuration, the **fapolicyd** service did not correctly identify users from a network source, which caused errors. This update fixes the security policy to allow the necessary communication. As a result, you can use **fapolicyd** with rules that require a network connection to identify users.

Jira:RHEL-21777

7.3. 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
```

Jira:RHELPLAN-137234[1]

7.4. SOFTWARE MANAGEMENT

dnf download --url correctly reports package URLs

Before this update, when you used the **dnf download --url** command to obtain a package URL, DNF incorrectly reported package addresses relative to the repository metadata location instead of relative to the **xml:base** attribute.

With this update, DNF considers the **xml:base** attribute when calculating package URLs. As a result, **dnf download --url** reports the correct package URL.

Jira:RHEL-71125[1]

7.5. SHELLS AND COMMAND-LINE TOOLS

/var/lib/tftpboot directory is created by default in Image Mode deployments

Previously, in Image Mode deployments, installing the **tftp-server** package did not create the **/var/lib/tftpboot** directory. This occurred because changes to the **/var** directory were not applied when additional packages were added to existing Image Mode deployments.

With this update, the **/var/lib/tftpboot** directory is automatically created in all Image Mode deployments.

Jira:RHEL-77491[1]

The top -u command now displays at least one process when you sort the processes by memory

Previously, when you executed the **top** command with the **-u *<user>*** parameter, where the ***user*** was different from the one running the command, all processes disappeared when the **M** key was pressed to sort the processes by memory. With this update, the **top** command displays at least one process when you sort the processes by memory.

NOTE

To preserve the position of the cursor, not all processes are displayed. You can scroll up through the results to display the remaining processes.

Jira:RHEL-46760

7.6. INFRASTRUCTURE SERVICES

The chronyc reload sources command now correctly handles hostname-specified sources

Previously, the **chronyc reload sources** command in **chronyd** incorrectly reloaded sources from the **sourcedir** directory specified in the **chrony.conf** file. This behavior caused the **chronyd** to duplicate sources when a hostname resolved to multiple IP addresses, resulting in an unexpected increase in the number of sources.

With this update, the **chronyc reload sources** command correctly handles sources specified with a hostname. As a result, reloading of sources does not change the number of used sources.

Jira:RHEL-95016

httpd works correctly if a DAV repository location is configured by using a regular expression match

Previously, when you configured a Distributed Authoring and Versioning (DAV) repository in the Apache HTTP Server by using a regular expression match, such as **LocationMatch**, the **mod_dav httpd** module was unable to determine the root of the repository from the path name. As a consequence, **httpd** did not handle requests from third-party providers, for example, Subversion's **mod_dav_svn** module.

This update introduces a new **DavBasePath** directive for the **httpd.conf** file, which allows you can specify the repository root path explicitly. For example:

```
<LocationMatch "^/repos/">
   DAV svn
   DavBasePath /repos
   SVNParentPath /var/www/svn
</LocationMatch>
```

As a result, **httpd** can correctly handle requests when you configure a DAV repository location by using a regular expression match.

Jira:RHEL-79948

httpd works correctly if a DAV repository location is configured by using a regular expression match

Previously, if a Distributed Authoring and Versioning (DAV) repository was configured in the Apache HTTP Server by using a regular expression match, such as **LocationMatch**, the **mod_dav** and **httpd** modules were unable to determine the root of the repository from the path name. As a consequence, **httpd** did not handle requests from third-party providers, such as Subversion's **mod_dav_svn** module.

With this update, you can specify the repository root path by using the new **DavBasePath** directive in the **httpd.conf** file. For example:

```
<LocationMatch "^/repos/">
   DAV svn
   DavBasePath /repos
   SVNParentPath /var/www/svn
</LocationMatch>
```

As a result, **httpd** handles requests correctly if a DAV repository location is configured by using a regular expression match.

Jira:RHEL-41069

The DBD::MySQL driver no longer fails to establish TLS-encrypted connections to MySQL 8 servers that have caching_sha2_password enabled

In previous releases, the **perl-DBD-MySQL** package was incorrectly linked against the **libmariadb** library. Consequently, Perl applications failed to establish a connection if all of the following conditions were met:

- The application connected to a MySQL 8 server.
- The **caching_sha2_password** option was enabled in the MySQL server configuration.
- The connection used the **DBI→connect with mysql_ssl=1** option.

In this update, the driver is linked against **libmysql-client**. As a result, Perl applications no longer fail to establish TLS-encrypted connections in the mentioned scenario.

Jira:RHEL-77083

7.7. NETWORKING

The custom iproute2 settings in /etc/iproute2/ works as expected

Previously, if you updated to RHEL 9.6, the **iproute2** package stored the default configuration in the **/usr/share/iproute2/** directory. Additionally, if you had a custom configuration in **/etc/iproute2/**, the update renamed these files and appended the **.rpmsave** suffix. As a consequence, the custom settings were no longer applied. If you update to the RHEL 9.7 version of the **iproute2** package, the installation script in the package no longer renames custom configuration files and, if it finds files with **.rpmsave** suffix in **/etc/iproute2/**, the script removes this suffix. As a result, custom settings work again as expected.

Note that the **iproute2** default settings remain in **/usr/share/iproute2/**.

Jira:RHEL-94662

The kernel no longer panics if you reduce the number of SR-IOV VFs at runtime

In previous releases, the Linux kernel could panic if all of the following conditions applied:

- The host has Input-Output Memory Management Unit (IOMMU) enabled.
- A network driver uses a page pool.
- You reduced the number of Single Root I/O Virtualization (SR-IOV) Virtual Functions (VFs) of the network interface that uses this driver.

With this update, the kernel tracks which DMA-mapped memory pages belong to a page pool. When a page pool is destroyed, for example by removing a VF, the memory pages are DMA-unmapped. This prevents attempts to unmap the memory pages after the VF has already been removed. As a result, the kernel no longer panics if you reduce the number of SR-IOV VFs at runtime.

Jira:RHEL-76845[1]

The xtables modules are now again marked as deprecated

Before this update, the **iptables**, **ip6tables**, **arptables**, **ebtables**, and **ip_set** driver were erroneously marked as unmaintained. As a consequence, RHEL logged an **Unmaintained driver is detected: *<driver>*** warning. With this release, the mentioned drivers have been marked again as deprecated. As a result, the system no longer reports the warning with the incorrect support status.

Jira:RHEL-81900

The xdp-loader features command now works as expected

The **xdp-loader** utility was compiled against the previous version of **libbpf**. As a consequence, **xdp-loader features** failed with an error:

> Cannot display features, because xdp-loader was compiled against an old version of libbpf without support for querying features.

The utility is now compiled against the correct **libbpf** version. As a result, the command now works as expected.

Jira:RHEL-3382

Mellanox ConnectX-5 adapter works in the DMFS mode

Previously, while using the Ethernet switch device driver model (**switchdev**) mode, the **mlx5** driver failed if configured in the device managed flow steering (**DMFS**) mode on the **ConnectX-5** adapter. Consequently, the following error message appeared:

> mlx5_core 0000:5e:00.0: mlx5_cmd_out_err:780:(pid 980895): DELETE_FLOW_TABLE_ENTRY(0x938) op_mod(0x0) failed, status bad resource(0x5), syndrome (0xabe70a), err(-22)

As a result, when you update the firmware version of the **ConnectX-5** adapter to 16.35.3006 or later, the error message will not appear.

Jira:RHEL-9897[1]

VMware vCenter can now correctly remove a SATA disk from a running RHEL VM

Previously, when using the VMWare vCenter interface to remove a SATA disk from a running RHEL 9 guest on the VMware ESXi hypervisor, the disk did not get removed fully. It stopped being functional and disappeared from the guest in the vCenter interface, but the SCSI interface still detected the disk as attached in the guest. With this update, the SCSI interface correctly displays the disk as detached.

Jira:RHEL-79914[1]

7.8. KERNEL

Updated the stalld scheduling policy regression to prevent performance degradation

Before this update, the Node Tuning Operator CI was broken because of a change in **stalld** scheduling policy., This change caused the service to revert to SCHED_OTHER instead of SCHED_FIFO after starting. Consequently, real-time workloads could experience performance degradation, and you could not merge PR. With this update, the **systemd** unit file sets **stalld** priority to 10, ensuring that **stalld** runs with SCHED_FIFO. This restores expected behavior and improves performance for real-time workloads.

Jira:RHEL-108827

osnoise/cpus allows setting a long comma-separated list of cpus

Before this update, you could not set a lengthy comma-separated list of cpus in osnoise/cpus because of an invalid argument error. This restriction impacted latency debugging and troubleshooting. With this release, you can input a long comma-separated list of cpus in osnoise/cpus to enhance RTLA latency debugging and troubleshooting.

Jira:RHEL-94317[1]

irqbalance service buffer overflow on aarch64 systems

Previously, the **irqbalance** service could crash due to a buffer overflow when running on specific aarch64 machines. As a consequence, latency-sensitive workloads might have experienced performance degradation because interrupts were not appropriately distributed across CPUs. With this update, the buffer overflow issue in the **irqbalance** service has been fixed.

As a result, the **irqbalance** service runs reliably, and interrupts are distributed as expected, improving performance for latency-sensitive workloads.

Jira:RHEL-89986

rtla timerlat does not reset osnoise stop tracing threshold during startup

Before this update, using the **rtla timerlat** multiple times without clearing the stop_tracing flags would leave/left **RTLA** in an inconsistent state. As a consequence, tracing did not stop correctly in case stop tracing was not requested via the -a, -T, or -i options. This led to inaccurate data being reported, since **RTLA** exited when it shouldn't have. With this update, **rtla-timerlat** resets stop tracing variables, preventing early exit, and as a result, program stability is improved.

Jira:RHEL-86051[1]

rtla timerlat now handles high-frequency sampling on systems with 100+ CPUs

Before this update, **rtla timerlat** could not process timerlat samples with 100us period or faster on systems with more than 100 CPUs due to insufficient **tracefs** buffer handling. As a consequence, samples were dropped and **timerlat** measurements became inaccurate, affecting real-time performance

analysis. With this release, **timerlat** samples are collected directly on measurement CPUs, eliminating buffer overflow issues. As a result, rtla timerlat provides accurate measurements on high-core-count systems, enabling reliable real-time performance analysis.

Jira:RHEL-77358[1]

7.9. FILE SYSTEMS AND STORAGE

multipathd can monitor devices with offline paths

Before this update, when a user created a multipath device while some paths to the device were in the offline state, the **multipathd** daemon did not monitor the device or its paths. Consequently, if paths failed, they were never restored, even if they became available again. With this update, the **multipathd** daemon monitors the multipath device and its offline paths. **multipathd** also adds the paths to the multipath device if they become online.

Jira:RHEL-82534

VDO driver no longer crashes due to null pointer dereference

Before this update, writing a mix of new and duplicate data to a VDO device under certain timing conditions left a dangling pointer. As a consequence, this caused a null pointer dereference and system crash. With this release, the dangling pointer issue is fixed. As a result, the VDO driver continues to run and saves user data.

Jira:RHEL-83857

The RHEL installation program removes corrupted LVM thin volumes

Previously, the presence of corrupted LVM thin volumes caused storage configuration errors, blocking the installation process. With this fix, the RHEL installation program now detects and removes broken thin volumes. As a result, users do not have to intervene in the installation process manually.

Jira:RHEL-8012

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]

7.10. HIGH AVAILABILITY AND CLUSTERS

pcs commands no longer fail due to improperly capitalized target-role values

Before this update, if a resource's **target-role** meta-attribute was set to a value that was not capitalized, such as **stopped** instead of **Stopped**, **pcs** failed to parse the cluster status. This parsing error caused **pcs status query resource** commands and commands for deleting resources, including **pcs resource delete**, to fail.

With this update, the cluster status parsing logic in **pcs** has been made more flexible.

As a result, **pcs** commands function correctly even when a resource has a **target-role** meta-attribute with an improperly capitalized value.

Jira:RHEL-92044

fence_ibm_powervs supports plain text token files

Before this update, the **fence_ibm_powervs** agent could only read authentication tokens from files that were formatted as JSON. It failed to read tokens from plain text files.

With this update, the file reading logic in the agent has been corrected.

As a result, the **fence_ibm_powervs** agent can use token files that are in either JSON or plain text format.

Jira:RHEL-88568

systemd resources with long start or stop times are handled correctly

Before this update, Pacemaker polled for the result of start and stop actions on **systemd** resources with a fixed timeout. If a resource took longer to start or stop than this timeout, Pacemaker incorrectly marked the resource as failed.

With this update, Pacemaker listens for DBus messages from **systemd** to be notified when a start or stop action completes.

As a result, Pacemaker correctly detects the status of long-running **systemd** services, and resources are no longer marked as failed due to a timeout.

Jira:RHEL-86143[1]

Pacemaker Remote nodes are no longer fenced unnecessarily when quorum is lost

Before this update, in certain cluster configurations, a Pacemaker Remote node could be fenced when its partition lost quorum, even if the resource managing that node could be safely restarted on a different, quorate node. This behavior caused unnecessary downtime for the services running on the Pacemaker Remote node.

With this update, a new cluster property, **fence-remote-without-quorum**, has been introduced to control this behavior.

As a result, with the default **fence-remote-without-quorum=false** setting, Pacemaker no longer fences a remote node if its managing resource can be recovered on a quorate node, thus improving service availability.

Jira:RHEL-84018[1]

fence_kubevirt powers off nodes instantly

Before this update, the **fence_kubevirt** agent performed a graceful shutdown of the node. This introduced a delay in the fencing process, as the node was not powered off immediately.

With this release, the agent has been modified to request an immediate, non-graceful shutdown.

As a result, when using the **fence_kubevirt** agent, nodes are instantly powered off.

Jira:RHEL-82193

fence_sbd is now more resilient to individual SBD device failures

Previously, the **fence_sbd** agent exited and failed its operation if one or more of its configured SBD devices failed an initial check. This prevented a fencing action from completing, even if other SBD devices were healthy.

With this update, the error handling in the agent is improved.

As a result, the **fence_sbd** agent logs an error for any failing SBD devices and continues the fencing operation with the remaining healthy devices. This increased the reliability of SBD fencing.

Jira:RHEL-13088[1]

7.11. COMPILERS AND DEVELOPMENT TOOLS

Improved support for recursive dlopen calls in audit modules in glibc

Previously, recursive **dlopen** calls from auditors could trigger an **r_state == RT_CONSISTENT** assertion failure in glibc's **dl-open.c**. As a consequence, applications exited unexpectedly when auditors were active. With this update, the dynamic linker reports consistency of its internal data structures earlier during an in-progress **dlopen** call. As a result, recursive **dlopen** operations for auditors are supported in more cases.

Jira:RHEL-47403

glibc: Application crash during early TLS allocation in audit mode

Previously, in audit mode, an internal data structure related to thread-local storage (TLS) management was allocated using the main **realloc** function before the main **malloc** was initialized during process startup. As a consequence, applications crashed when **realloc** was called on memory that was not allocated by **malloc**.

With this update, the dynamic linker uses a stub or minimal implementation of **malloc** and **realloc** until the startup process is complete. The applications no longer crash during early TLS allocation.

Jira:RHEL-71922[1]

glibc: ctype.h macros caused segmentation faults in multithreaded programs with multiple libc.so

Previously, the internal state for **<ctype.h>** in secondary C library copies created by audit or with **dlmopen** failed to initialize for threads created with **pthread_create**. As a consequence, using **<ctype.h>** functionality, either directly or indirectly, in secondary threads and namespaces resulted in program crashes.

With this update, the internal state for **<ctype.h>** is initialized to refer to the **C** locale for secondary threads and namespaces. As a result, using functionality from **<ctype.h>** in these scenarios no longer causes crashes.

Jira:RHEL-72017

glibc audit logging provides complete object life cycle tracking

Previously, the glibc dynamic linker called **la_objclose** for the proxy **ld.so** link map in a secondary namespace without a preceding **la_objopen**, which resulted in incomplete object life cycle reporting for tools that rely on **la_objopen** to track shared objects.

Auditing tools that rely on **la_objopen** to establish tracking failed to monitor proxy link maps reliably, resulting in gaps in visibility and possible misinterpretation of unload events.

With this update, the glibc dynamic linker generates **la_objopen** for the applicable link maps, including the proxy **ld.so** in secondary namespaces, ensuring a consistent sequence for the auditing interface.

As a result, auditors can track proxy link maps throughout their life cycle with consistent **la_objopen** and **la_objclose** pairs, improving the reliability of audit tooling and diagnostics.

Jira:RHEL-49549

Certain programs no longer crash when running glibc in auditing mode

Before this update, the **glibc** dynamic linker in **LD_AUDIT** mode could allocate internal data structures by using the main **calloc** function before the linker initialized the main **malloc** subsystem. As a consequence, the process could terminate unexpectedly in the **calloc** function when the program started. This update rearranges the process startup sequence. Consequently, the **calloc** memory allocation occurs before the switch to the main **malloc** function by using the internal **malloc** implementation, which is used during the startup. As a result, programs no longer crash during startup in the **calloc** function when running if the dynamic linker uses the auditing mode.

Jira:RHEL-48820[1]

stdio flushing issues fixed in glibc

Before this update, specific stdio streams in glibc could fail during fclose when attempting to seek back to the correct position after buffered reads, returning EINVAL instead of the expected ESPIPE for non-seekable inputs. As a consequence, applications using fclose on pipes or other non-seekable descriptors might encounter unexpected errors, causing I/O cleanup to fail and leading to inconsistent file positioning behavior.

With this update, glibc synthesizes an ESPIPE error when lseek returns **0** after bytes are read, ensuring fclose ignores the non-seekable condition as intended, and supporting test infrastructure changes (for example, xdup) to validate the behavior. As a result, fclose and related stdio operations now behave consistently for non-seekable streams, reducing error conditions and improving reliability in applications that rely on buffered I/O.

Jira:RHEL-68805[1]

Applications no longer deadlock when invoking popen and fork in parallel

Before this update, when multi-threaded applications invoked **popen** and **fork** in separate threads and the **fork** occurred when **popen** held an internal lock, the child process could inherit the locked state and deadlock if it called **popen** again.

With this update, glibc releases the relevant lock state across **fork**, ensuring that subsequent **popen** calls proceed without blocking. As a result, **popen** no longer deadlocks after a multi-threaded **fork** call, improving process reliability and input-output behavior for supported architectures.

Jira:RHEL-59712[1]

Golist command-line parser fix in go-rpm-macros

Before this update, Golist handled certain files incorrectly due to a replacement of the command-line parser. As a consequence, some programs failed to build. With this update, the original command-line parser restored the original parser into Golist.

As a result, Golist processes all required files correctly and programs are built as expected.

Jira:RHEL-7366

7.12. IDENTITY MANAGEMENT

ipa-cacert-manage install now permits duplicate CA subjects

Previously, attempting to add a CA certificate with an identical subject but a different private key using **ipa-cacert-manage install** failed with the message **subject public key info mismatch**, as IdM prohibited duplicate subjects.

This update relaxes that restriction, allowing **ipa-cacert-manage install** to accept duplicate CA subjects. However, the following limitations remain:

- Certificates cannot be added with different trust flags.
- The CAs must share the same nickname.
- An Authority Key Identifier (AKI) extension is mandatory for all CAs. Its absence leads to an unexpected chain of trust behavior.

Jira:RHEL-30658[1]

Newly created user password policies are displayed correctly

Before this update, the **cosAttribute** attribute in the Class of Service (CoS) template had the **operational** modifier instead of **operational-default**. As a consequence, when both subtree and user password policies existed, the **pwdpolicysubentry** attribute pointed to the subtree password policy instead of the user password policy. With this release, the CoS template uses the **operational-default** modifier. As a result, the user policy is displayed correctly.

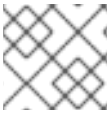

NOTE

This issue affected only displaying the policies, not the actual password policy logic.

Jira:RHEL-109892[1]

The RootDN Access Control plugin with wildcards for IP addresses no longer fails

Before this update, if you tried to set IP addresses with wildcards for the RootDN Access Control plugin configuration, the attempt failed with the **Invalid IP address** error. With this release, the validation function was updated. As a result, the attempt to set values with wildcards no longer fails.

Jira:RHEL-109889[1]

The Databases menu opens as expected in the Directory Server web console

Before this update, you could not open the **Databases** menu in the Directory Server web console if the database name that you created had an incorrect suffix syntax, for example, the name included **dc=**. With this update, Directory Server uses a rollback functionality when mapping tree creation fails during backend creation to prevent orphaned backends. As a result, the **Databases** menu opens as expected.

Jira:RHEL-109885[1]

Directory Server no longer fails when adding nsslapd-referral

Before this update, when you tried to configure Directory Server to use a referral, the incorrect handling of the paged search result caused the server failure.

With this update, If the search result code is **LDAP_REFERRAL**, the paged result search returns the correct value and the server no longer fails.

Jira:RHEL-107585[1]

The Directory Server monitoring information is available as expected when NDN cache is disabled

Before this update, when the Normalized DN (NDN) cache was disabled, the **dsconf *<instance_name>* monitor dbmon** command failed with an error because of improper handling of the backend get-tree command failures. This release adds a rollback functionality to prevent orphaned backends when the tree creation fails during a backend creation. As a result, Directory Server monitoring information is returned as expected.

Jira:RHEL-107005[1]

Directory Server correctly displays the number of child entries under a specific node

Before this update, the **numSubordinates** and **numTombstoneSubordinates** attributes were wrongly computed during import. Consequently, when you compared the number of child entries under a specific node, the wrong values were displayed.

With this update, Directory Server computes **numSubordinates** and **numTombstoneSubordinates** correctly.

Jira:RHEL-104593[1]

The Directory Server web console now shows the server version

Before this update, the web console did not display the server version in the **Server Settings>General Settings**. With this update, the server version is displayed correctly.

Jira:RHEL-104591[1]

Directory Server no longer fails during NDN cache operations

Before this update, the **arc-swap** library, which was used in the Rust dependency of **389-ds-base**, could cause a failure in Directory Server during NDN cache operations. With this release, Directory Server uses an updated version of Rust dependency (concread) 0.5.7 that does not contain the **arc-swap** library. As a result, Directory Server no longer fails.

Jira:RHEL-95444[1]

Directory Server correctly displays membership in nested groups

Before this update, Directory Server displayed an incorrect value of the **memberOf** attribute in that entry under the following conditions:

- An entry was a member of groups that had multiple nested levels
- Groups were part of other different groups that had multiple paths in the membership relations.

With this update, the **memberOf** distinguished name (DN) value is added systematically, and the entry membership in groups is displayed correctly.

Jira:RHEL-89753[1]

389-ds-base no longer fails during the LMDB offline import

Before this update, a race condition occurred when a worker thread read an entry before another process finished writing the entry. As a result, offline import on an instance with the Lightning Memory-Mapped Database Manager (LMDB) backend caused a segmentation fault.

With this update, Directory Server ensures thread-safe access by locking the worker queue before writing entries, and the server no longer fails during the LMDB offline import.

Jira:RHEL-89745[1]

dsconf correctly returns replication monitoring information

Before this update, if a supplier was configured with a replica starting with **0**, such as **010** or **020**, the **dsconf <instance_name> replication monitor** command failed to retrieve information about time of a delay or the replication status.

With this update, non-significant zeros (**0**) at the beginning of replica ID are ignored while processing the replica ID within the replica update vector (RUV). As a result, **dsconf <instance_name> replication monitor** provides the expected information.

Jira:RHEL-89736[1]

ipa-healthcheck now ignores the replica busy condition

Before this update, in a topology with more than two suppliers, the **ipa-healthcheck** tool reported an error about replication agreement status when a supplier was receiving updates from another node. It is a standard replication situation and, with this release, **ipa-healthcheck** no longer reports an error when replicas are busy.

Jira:RHEL-79673

Directory Server starts correctly in the read-only mode

Before this update, Directory Server did not start if you configured the read-only mode. With this update, the **nsslapd-readonly** attribute is processed correctly, and the server starts in the read-only mode as expected.

Jira:RHEL-61347

Obsolete /var/log/tallylog log file creation removed

Before this update, an outdated configuration in the **pam.conf** file caused the creation of the **/var/log/tallylog** file. Since the system now uses **pam_faillock**, which replaced the obsolete **pam_tally**, the **/var/log/tallylog** file is no longer necessary.

With this update, **pam.conf** was updated to remove the instructions for the obsolete log file creation.

Jira:RHEL-15324

7.13. DESKTOP

Default GDM session definitions no longer override custom definitions

Before this update, GNOME Display Manager (GDM) sessions at **/usr/** directories had higher precedence than the ones at **/etc/**. As a consequence, custom session definitions at **/usr/** would override the ones at **/etc/**. With this release, sessions at **/etc/** have higher precedence. As a result, the custom definitions precedence works correctly as defined in GDM Session Configuration.

Jira:RHEL-95837

7.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES

encryption_key is no longer masked

Before this update, the **encryption_key** parameter was incorrectly marked as **no_log**. This caused the key file path to be replaced by a placeholder string, preventing disk encryption from working. With this update, the **encryption_key** parameter is no longer marked with the **no_log** flag, and you can now perform disk encryption using a key file successfully.

Jira:RHEL-104676[1]

RAID now reports clear errors for invalid or unsupported configurations

Before this update, invalid RAID levels or insufficient disks could be specified without raising clear errors. This resulted in failed or inconsistent array creation. As a consequence, the error messages were unclear, and RAID setup was less reliable. With this release, RAID parameters are validated before array creation, and a minimum disk count is enforced. As a result, clear errors are raised, and attempts to create a RAID with inadequate disks are blocked.

This fix also removes the deprecated **process_device_numbers** helper and uses **unify_raid_level** instead. In addition, failure tests for invalid RAID level and insufficient disks scenarios are also added.

Jira:RHEL-104891[1]

LVM RAID now supports encrypted and partitioned devices

Before this update, the LVM RAID code assumed that disks specified in **raid_disks** were the parent devices of the PVs for all LVM RAID setups. This was not applicable for encrypted or partitioned devices. As a consequence, errors occurred when encrypted LUKS layers added an extra storage layer, or when direct partitions were used without a parent device. With this release, PV resolution in LVM RAID is improved to support encrypted and partitioned devices. As a result, you can now specify the PV partition instead of the underlying disk.

This fix also adds error handling for missing or invalid RAID disk entries and introduces corresponding tests to ensure stability.

Jira:RHEL-95885[1]

Minor volume size mismatch no longer cause incorrect role reporting

Before this update, when creating or resizing volumes, the system allowed up to a 2% difference between the requested size and the actual size. This adjustment made the volume fit into the available pool free space. As a consequence, the sizes did not match when the role was run again, causing the role to incorrectly assume that something had changed. With this release, small size differences no longer cause the role to misinterpret changes. As a result the role now reports the correct state.

Jira:RHEL-82825[1]

The postfix RHEL system role auto-detects if an IPv6 interface is disabled

The default **postfix** configuration uses the **inet_interfaces = localhost** setting which tells **postfix** to listen on all interfaces resolving to **localhost** including both IPv4 and IPv6 interfaces. Before this update, a problem occurred if IPv6 was disabled on the host. In this situation, the **postfix** role and its command-line tools, such as **postconf**, returned an error. The entire role failed. With this release, the role determines if IPv6 is disabled. If so, then it sets **inet_protocols = ipv4** so that **postfix** only uses the IPv4 interface. As a result, the **postfix** role works even when IPv6 is disabled.

Jira:RHEL-103889[1]

The timesync RHEL system role no longer removes the OPTIONS="-F 2" default setting from /etc/sysconfig/chronyd

Before this update, the **timesync** system role replaced the default **OPTIONS=** setting for the **chronyd** service with **""**. As a consequence, this removed the default **OPTIONS="-F 2"** setting which weakened the security of **chronyd**. With this release, **-F 2** is added as the default setting for **OPTIONS**, and the user can override or extend this setting. As a result, the **timesync** role now applies the correct security settings while still allowing user customization.

Jira:RHEL-88299[1]

Improved removal of kernel options with values in rhel-system-roles

Previously, kernel boot options specified as key=value could not be removed when users provided only the key, resulting in persistent unwanted boot parameters and inconsistent management of kernel options by name. With this update, the regular expression in the **mod_boot_args** function was updated to match and remove kernel options with values correctly, and automated tests were added to verify correct behavior.

As a result, kernel options can now be reliably removed by name, even when set as key=value, ensuring accurate configuration and improved system management.

Jira:RHEL-101678[1]

GSSAPIIndicators added to sshd role

A new configuration option **GSSAPIIndicators** for setting Generic Security Services Application Programming Interface (GSS-API) was added to RHEL 10. This update adds the **GSSAPIIndicators** configuration option to the **sshd** RHEL system role. As a result, you can configure **GSSAPIIndicators** on RHEL 10 systems by using RHEL system roles.

Jira:RHEL-107049[1]

bootloader role rejects boolean or null type values

Before this update, the user could specify values such as **value: on** or **value: yes** expecting that these would be converted to strings **"on"** or **"yes"**. But instead, YAML treats these as YAML bool type and writes them as the string **"True"**. Consequently, users who were unaware of YAML boolean handling could not set values such as **"on"** or **"off"**. With this update, the **bootloader** RHEL system role rejects any value of boolean or **null** type. As a result, users must enter such YAML boolean type values as quoted strings to write them to the bootloader configuration. The readme is updated with this information.

Jira:RHEL-107015[1]

sudo role no longer hangs when parsing Alias values

Before this update, the regex in the **sudo** RHEL system role was not taking into consideration that Alias values, such as **Cmnd_Alias**, do not have to have spaces on either side of the equal sign **=**. Consequently, the regex never terminated, and the role appeared to hang. With this update, the role ensures that the regex complies with the eBNF definition of the field from the **sudoers** file specification. As a result, the Alias values are parsed correctly with and without spaces around **=**.

Jira:RHEL-106733[1]

Specifying multiple users no longer causes resources to be associated with wrong user

Before this update, user data contamination occurred due to mixing facts and variables for the **__podman_user** and **__podman_user_home_dir** variable values when managing multiple users. As a consequence, user data was mixed between multiple users, causing incorrect configuration files to be used for each user. With this release, user data separation is maintained by avoiding the mixing of facts and variables for **__podman_user** and **__podman_user_home_dir**. As a result, user data is isolated for multiple users, improving resource management consistency.

Jira:RHEL-105095[1]

selinux role no longer produces error due to undefined tempdir path in Ansible check mode

Before this update, the **tempdir** path was not defined in Ansible check mode, and the **__selinux_item.path** could be undefined. Consequently, when running in check mode, the **selinux** RHEL system role produced an error that various variables are undefined. With this update, the role skips tasks that require the **tempdir.path** to be defined, and can handle cases where variables are undefined. As a result, the role works correctly in check mode.

Jira:RHEL-103575[1]

Ensures /var/lib/pcsd directory is available when needed by the ha_cluster RHEL system role

Before this update, the **/var/lib/pcsd** directory was created during the installation of **pcs**, but newer versions rely on the **systemd** service to create this directory when the **pcsd** service starts. As a result, the directory might not exist at the time the role attempts to access it, causing errors or failures in execution.

With this update, the role explicitly ensures that the **/var/lib/pcsd** directory exists before using it. As a result, it prevents runtime issues due to the missing directory and improving the reliability of role execution.

Jira:RHEL-101663[1]

Using the redhat.rhel_system_roles collection no longer displays a warning about an incompatible Ansible version

Before this update, the **redhat.rhel_system_roles** collection specified **{{requires_ansible: ">=2.15.0"}}** in the **meta/runtime.yml** file, but RHEL 9 contains **ansible-core** 2.14. As a consequence, if you used the collection in a playbook, Ansible displayed a **Collection redhat.rhel_system_roles does not support Ansible version 2.14.x** warning. This update changes the **meta/runtime.yml** file to use **{{requires_ansible: ">=2.14.0"}}**. As a result, the collection no longer displays the warning.

Jira:RHEL-94444[1]

selinux role persistently sets kernel SELinux parameters

Before this update, the **selinux** RHEL system role did not set the kernel SELinux parameter when changing the SELinux state to and from disabled. As a consequence, the SELinux state change was not persistent upon reboot. This update ensures that the kernel SELinux parameter is correctly set when the role changes SELinux state to and from disabled. As a result, the SELinux state change to and from disabled is persistent upon reboot.

Jira:RHEL-93296[1]

The systemd role uses file basename to construct the path to the destination

Before this update, if a user specified a file or a template source within a nested directory, the **systemd** RHEL system role used the whole path instead of the basename for the destination file. As a consequence, files and templates were placed in the same directory structure on the destination, which **systemd** does not support. With this release, the role uses basenames for destination files in nested directories. As a result, users can use nested directories with the role.

Jira:RHEL-88772

Introducing flexibility for package installation in ad_integration role

Previously, the **ad_integration** role always attempted to install the required packages, for example, **realmd**, **sssd-ad**, **adcli**, and many more that are listed in **__ad_integration_packages**. In environments where external systems handled package management, for example, via configuration management outside of this role, pre-baked images, or immutable systems, this step was redundant and undesirable.

With this update, users can now manage package installations through other means and only want this role to join a domain, offering them flexibility. The notable enhancements are:

- New Variable: Introduced a new boolean variable **ad_integration_manage_packages** to control whether the role installs packages.
- Default Value: The default value is set to **true** in **defaults/main.yml** to ensure backward compatibility. Existing playbooks using this role will continue to function as before without modification.
- Conditional Task: Added a **when: ad_integration_manage_packages | bool** condition to the "Ensure required packages are installed" task in **tasks/main.yml**. The task will now only run if the flag is **true** (the default).
- Documentation: Updated **README.md** to include the new **ad_integration_manage_packages** variable, explaining its purpose and default value.

Jira:RHEL-88314[1]

The qdevice daemon now restarts automatically after certificate changes

Previously, after updating the TLS certificates used for communication between the quorum device daemon (**qnetd**) and the cluster nodes (**qdevice**), the **qdevice** daemon was not automatically restarted. The daemon would continue to use the old certificates, causing communication with the quorum device to fail.

With this update, the **qdevice** daemon on cluster nodes automatically restarts after its certificates are changed. This ensures that the new certificates are loaded immediately and that communication with the quorum device is maintained.

Jira:RHEL-88251[1]

The ha_cluster RHEL System Role now works with a system-wide HTTP proxy configured

Previously, when a system-wide HTTP proxy was configured, the **ha_cluster** RHEL System Role would incorrectly attempt to use the proxy for local communication with the **pcsd** daemon via a unix socket. This caused the role to fail.

With this release, the role has been modified to explicitly disable proxy usage for local **pcsd** communication.

As a result, the **ha_cluster** RHEL System Role works as expected on systems with a system-wide HTTP proxy defined.

Jira:RHEL-88241[1]

The network RHEL system role no longer shows errors due to incorrect routing rule validation

Before this update, the validation part in the **network** RHEL system role incorrectly checked for routing rule attributes at the top-level **NM** module instead of the **NM.IPRoutingRule** class. This caused validation failures and the role displayed errors. With this update, the role uses the API correctly and no longer shows incorrect validation errors.

Jira:RHEL-85872

Boolean option values are correctly rendered in TOML files

Previously, the boolean options were mishandled because the formatter code did not convert the boolean values to the correct string representation. With this fix, boolean values are properly converted to lowercase strings, ensuring correct rendering and handling in TOML files.

Jira:RHEL-85702

Boolean options are correctly written and handled in TOML files

Before this update, boolean options were not correctly handled because the code that formats into TOML format did not convert boolean values to the correct string representation.

With this update, we convert boolean options to string and then to lower case, which is the correct TOML boolean format. This ensures that TOML files correctly write and handle boolean options.

Jira:RHEL-84940

The podman RHEL system role does not report changed: true when managing authentication and configuration files

Before this update, the **podman** RHEL system role changed the parent path mode every time it ran if it managed both authentication and configuration files because it used two different modes for the common parent path for various configuration and authentication files.

With this fix, the role does not report **changed: true** unnecessarily because it uses a consistent mode for the parent path.

Jira:RHEL-84920

Podman role does not fail with UNREACHABLE error

Previously, the **podman** role did not wait enough for the user state to be in **closing** status when disabling linger for non-root users. The **podman** role then restarted **systemd-logind** to force it to

cancel. On some systems, this started a timer that killed the session for root, causing the **sshd** session to terminate and the Ansible play to fail with **UNREACHABLE** error.

With this fix, the system now waits much longer for users to be in the **closing** state, and only restarts **logind** if absolutely necessary. As a result, the role does not fail with **UNREACHABLE** error when removing resources.

Jira:RHEL-84910

The network RHEL system role now uses a more robust interface identification method

Before this update, when both an interface name and a MAC address were provided for a network interface, the validation process performed two separate lookups: one using the interface name and another using the MAC address. This could lead to validation failures because a lookup by MAC address might match the interface's current MAC address rather than its permanent hardware MAC address.

With this update, the validation logic has been improved. The network role now uses the interface name as the only identifier to look up the network device. It then retrieves the MAC address associated with that interface and compares it to the user-provided MAC address for validation. This approach is more reliable, because interface names are unique kernel identifiers, preventing mismatches caused by temporary MAC address changes.

Jira:RHEL-84362

The systemd role unmasks and starts units in a single run

Before this update, the **systemd** RHEL system role failed to enable and start services when units were masked because the role could not unmask the units first. As a result, users had to run the role twice. With this release, the **systemd** role correctly unmasks and starts services, eliminating the need for double runs.

Jira:RHEL-81755

7.15. VIRTUALIZATION

Local kdump no longer fails on virtual machines with AMD SEV-SNP

Before this update, local kdump failed on RHEL 10 virtual machines (VMs) that used the AMD Secure Encrypted Virtualization (SEV) with the Secure Nested Paging (SNP) feature. As a consequence, you could not capture kernel crash dumps on VMs with AMD SEV-SNP enabled.

With this release, the underlying code has been fixed. As a result, local kdump no longer fails on VMs with AMD SEV-SNP.

Jira:RHEL-10019[1]

The --migrate-disks-detect-zeroes option no longer fails for VM migration

Before this update, when migrating virtual machines (VMs) on RHEL 10, the **--migrate-disks-detect-zeroes** option might not have worked, and the migration might have proceeded without zeroed block detection on the specified disk. This problem was caused by a bug in QEMU where mirroring jobs relied on punching holes, resulting in a sparse destination file.

With this release, QEMU has been fixed to preserve sparseness if the destination system reports that it reads all zeroes, and when no extra effort is made to further sparsify the image. As a result, the **--migrate-disks-detect-zeroes** option works as expected for VM migration.

Jira:RHEL-82906

VMs sending misaligned discard I/O requests no longer pause when **discard_granularity** is not configured

Before this update, the host kernel failed misaligned discard I/O requests and QEMU used the **werror=policy** parameter to respond to such failures. When **werror** was set to **stop**: **werror=stop**, a failed discard request caused the virtual machine (VM) to pause. As a consequence, it was not possible to correct this situation and resume the VM again.

With this release, QEMU has been updated to silently ignore misaligned discard I/O requests, so that guests without a correct discard_granularity value do not pause. As a result, VMs sending discard I/O requests no longer pause when **discard_granularity** is not configured. However, it is still preferable to configure the **discard_granularity** value, so that discard requests have their intended effect instead of being ignored when misaligned.

Jira:RHEL-86032[1]

virtiofsd no longer crashes when accessing shared directories with many open files

Before this update, when accessing a **virtiofs** shared directory with a large number of open files from a virtual machine (VM), the operation might have failed with the following error: **Too many open files**, and the **virtiofsd** process crashed.

With this release, the underlying code has been fixed. As a result, accessing a **virtiofs** shared directory with a large number of open files from a VM might still result in an error in the VM, but the **virtiofsd** process no longer crashes, keeping the **virtiofs** shared directory accessible in the VM.

Jira:RHEL-87161[1]

Customizing RHEL 9 guests on ESXi no longer causes networking problems

Previously, customizing a RHEL 9 guest operating system in the VMware ESXi hypervisor did not work correctly with NetworkManager key files. As a consequence, if the guest was using such a key file, it had incorrect network settings, such as the IP address or the gateway. This problem has now been fixed, and NetworkManager key files no longer cause networking issues in the described scenario.

Jira:RHELPLAN-106947[1]

The installation program 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 installation program 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 installation program 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 installation program shows the expected system disk to install RHEL in the VM.

Jira:RHELPLAN-66975[1]

Windows guests boot more reliably 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 previously failed to boot. This occurred on hosts that use AMD EPYC series CPUs. Now, the underlying code has been fixed and VMs boot as expected in the described circumstances.

Jira:RHELPLAN-147926[1]

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.

Jira:RHELPLAN-139536[1]

virtiofs devices can now 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.

Jira:RHELPLAN-119912[1]

blob resources now work correctly for virtio-gpu on IBM Z

Previously, the **virtio-gpu** device was incompatible with **blob** memory resources on IBM Z systems. As a consequence, if you configured a virtual machine (VM) with **virtio-gpu** on an IBM Z host to use **blob** resources, the VM did not have any graphical output.

With this update, **virtio** devices have an optional **blob** attribute. Setting **blob** to **on** enables the use of **blob** resources in the device. This prevents the described problem in **virtio-gpu** devices, and can also accelerate the display path by reducing or eliminating copying of pixel data between the guest and host. Note that **blob** resource support requires QEMU version 6.1 or later.

Jira:RHEL-7135

Reinstalling virtio-win drivers no longer causes DNS configuration to reset on the guest

In virtual machines (VMs) that use a Windows guest operating system, reinstalling or upgrading **virtio-win** drivers for the network interface card (NIC) previously caused DNS settings in the guest to reset. As a consequence, your Windows guest in some cases lost network connectivity.

With this update, the described problem has been fixed. As a result, if you reinstall or upgrade from the latest version of **virtio-win**, the problem no longer occurs. Note, however, that upgrading from a prior version of **virtio-win** will not fix the problem, and DNS resets might still occur in your Windows guests.

Jira:RHEL-1860[1]

VNC viewer correctly initializes a VM display after live migration of ramfb

This update enhances the **ramfb** framebuffer device, which you can configure as a primary display for a virtual machine (VM). Previously, **ramfb** was unable to migrate, which resulted in VMs that use **ramfb** showing a blank screen after live migration. Now, **ramfb** is compatible with live migration. As a result, you see the VM desktop display when the migration completes.

Jira:RHEL-7478

7.16. RHEL IN CLOUD ENVIRONMENTS

Nested VM with KVM virtualization and OVMF now boots successfully on Azure or Hyper-V when using an AMD EPYC processor

Previously, a nested virtual machine (VM) with Open Virtual Machine Firmware (OVMF) failed to boot when run on a RHEL VM with KVM virtualization enabled on Microsoft Azure or Hyper-V that used an AMD EPYC processor. The VM failed to boot up with following log message:

```
Code=qemu-kvm: ../hw/core/cpu-sysemu.c:76  Aborted (core dumped) .
```

With this update, the problem has been fixed, and the nested VM boots as expected in the described circumstances.

Jira:RHEL-29919[1]

7.17. SUPPORTABILITY

The **coredump** plugin now correctly limits the number of collected **coredump** files

Previously, the **coredump** plugin collected **coredumpctl dump** outputs, which could lead to unnecessary large archives. With this update, the plugin defaults to collecting the three most recent **coredump** files. Additionally, the plugin continues to provide summary information from **coredumpctl info** and includes symlinks to help map collected dumps to their respective metadata entries.

Users can further filter collected dumps using the **executable** option, which accepts a case-insensitive Python regular expression applied to the EXE field of **coredumpctl list**. You can further use the **dumps** option to limit the number of last coredumps.

Jira:RHEL-62972[1]

Plugin option overrides in **sos report** no longer disable unrelated options configured in **/etc/sos/sos.conf** or a preset

Previously, when executing the **sos report** command with a **-k** option specifying a particular plugin setting , the **sos** utility would incorrectly ignore other valid plugin options defined in **/etc/sos/sos.conf** or in a preset. This led to scenarios where global settings or user-defined presets, were silently disabled despite being correctly configured in the **[plugin_options]** section of the configuration file or in a preset.

This behavior affected customers attempting to collect full System Activity Reporter (SAR) data as outlined in Red Hat Knowledgebase Solution 1418303 . When any **-k** option was used at runtime, the **sar.all_sar** setting reverted to **off**, resulting in incomplete data collection.

With this update, the **sos** tool now correctly merges options provided via the **-k** flag with those defined in the configuration file, ensuring that unrelated plugin options are preserved and applied as expected. This fix restores consistency and ensures comprehensive SAR data collection when configured.

Jira:RHEL-67097[1]

sos-audit package now includes required GPLv2 LICENSE file

Previously, while the **sos-audit** package was always part of the **sos** project and built from the same SRPM containing the license, the resulting **sos-audit** RPM package could be installed separately from

the main **sos** RPM. This meant users installing only the **sos-audit** subpackage would not find the license readily available. This omission affected all versions of **sos-audit** up to the current release across RHEL 8 and RHEL 9.

With this update, the **sos-audit** package now correctly includes the GPLv2 **LICENSE** file.

Jira:RHEL-73028

iscsi plugin no longer collects plain-text CHAP credentials in sosreport

Previously, the **iscsi** plugin in **sos** collected sensitive CHAP authentication credentials in **iscsi** configuration files in plain text when generating a report that posed a security risk. With this update, the **iscsi** plugin has been modified to obscure sensitive fields, ensuring that CHAP usernames and passwords are redacted or excluded from the collected output.

Jira:RHEL-81187[1]

THP plugin now collects complete configuration to accurately reflect Transparent Huge Pages state

Previously, the memory plugin of **sos** collected only the **enabled** file from **/sys/kernel/mm/transparent_hugepage/** to determine the state of Transparent Huge Pages (THP). However, recent kernel behavior changes have made this approach insufficient. For instance, it is possible for **enabled** to be set to **[never]** while **shmem_enabled** is set to **[always]**, resulting in THP being active for shared memory segments despite appearing disabled.

With this update, the THP plugin now collects all relevant files under **/sys/kernel/mm/transparent_hugepage/**, providing a complete and accurate view of how and where THP is enabled.

Jira:RHEL-81634[1]

per-user SSH configuration is now disabled by default

Previously, the **ssh** plugin in **sos** collected detailed information from all local user **.ssh** directories by default. This resulted in significantly prolonged execution time, especially in environments with a large number of local users. With this update, the **ssh** plugin no longer collects per-user **.ssh** configuration data by default. To capture user configurations, enable it explicitly by setting **ssh.userconfs=on**.

Jira:RHEL-84078

sos collect command in the sos 4.10 version no longer produces xz/bz2 tar archive

Before this update, the **sos collect** command returned a compressed tar archive like **tar.xz** or **tar.bz2**. With this release, the **sos collect** now produces uncompressed **tar** archives instead of compressed ones, saving time and resources.

Jira:RHELDOCS-21013[1]

7.18. CONTAINERS

Event logs from podman events command are now available

Previously, an error in the **journald** driver prevented the preservation of network event attributes, so these events were not included in logs. With this update, **podman events** now displays **network create** and **network rm** events.

Jira:RHEL-110317

Parent directories can be created now for the mount targets with mode 0755

In this update, build failures were occurring due to modifications in the handling of **--mount** parameter permissions in **quay.io/buildah/stable:v1 v1.41.3**. Previously, specifying UID as an argument resulted in incorrect permissions for the secret. Consequently, users were unable to access build secrets due to incorrect permissions after the **buildah** update.

With this release, Buildah has updated secret permissions for Buildah v1.41.3, using **secret-permissions** instead of **mount**. As a result, Buildah now sets the expected permissions for secrets correctly when using the UID argument in the **--mount** parameter, resolving mount failures.

Jira:RHEL-115166

7.19. RHEL LIGHTSPEED

Command-line assistant shows a meaningful error message when you try to delete a non-existent chat history

Before this update, users could delete a non-existent chat history without receiving an error message. This enhancement implements an error message for such cases.

Jira:RHELDOCS-21314[1]

Adding a description to an unnamed chat triggers a warning

Before this update, if you added a description to a chat without specifying a name for the chat, there was no error message displayed, nor was the chat with your custom description. With this update, the command-line assistant displays a warning in such cases.

Jira:RHELDOCS-21316[1]

c history shows complete history by default

Before this update, running the **c history** command without any options returned no history, confusing users. With this update, the default option for **--all** has been added. As a result, you can easily view all history with the single command: **c history**.

Jira:RHELDOCS-21317[1]

Command-line assistant no longer displays errors for invalid queries

Before this update, an incorrect data structure for terminal output in response led to unprocessable error messages for user queries. With this enhancement, the chat interface's terminal output structure has been actively addressed, preventing the command-line assistant from displaying errors for invalid query requests, thereby enhancing your user experience.

Jira:RHELDOCS-21318[1]

Interactive shell starts correctly after a terminal restart

Before this update, the user's **.bashrc** file did not include a reference to the **.bashrc.d** directory, preventing the **source** command from locating the CLA integration script. As a consequence, users could not access an interactive shell. With this update, a check has been added to ensure that the files necessary for shell integration are loaded. As a result, the interactive shell starts upon terminal restart.

Jira:RHELDOCS-21319[1]

Backend timeout works correctly in query.py

Before this update, extending the backend timeout in the **query.py** script did not work correctly. The script continued to generate timeout messages every 30 seconds because an internal timeout remained set at 30 seconds by default. With this enhancement, you can extend the backend timeout to any value that suits you by configuring this in the **/etc/xdg/command-line-assistant/config.toml** file, improving your response time.

Jira:RHELDOCS-21320[1]

cla chat displays help when run without arguments

Before this update, using **cla chat** without providing additional input caused user confusion, as they expected interactive AI assistance but received no response. With this update, when you use **cla chat** without arguments, the command-line assistant provides help and indicates additional input, improving your user experience with CLA's interactive mode.

Jira:RHELDOCS-21322[1]

CHAPTER 8. 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.

8.1. INSTALLER AND IMAGE CREATION

NVMe over TCP for RHEL installation is now available as a Technology Preview

With this Technology Preview, you can now use NVMe over TCP volumes to install RHEL after configuring the firmware. While adding disks from the Installation Destination screen, you can select the NVMe namespaces under the NVMe Fabrics Devices section.

Jira:RHEL-10216[1]

Installation of bootable OSTree native containers is now available as a Technology Preview

The **ostreecontainer** Kickstart command is now available in Anaconda as a Technology Preview. You can use this command to install the operating system from an OSTree commit encapsulated in an OCI image. When performing Kickstart installations, the following commands are available together with **ostreecontainer**:

- graphical, text, or cmdline
- ostreecontainer
- clearpart, zerombr
- autopart
- part
- logvol, volgroup
- reboot and shutdown
- lang
- rootpw
- sshkey
- bootloader - Available only with the **--append** optional parameter.
- user

When you specify a group within the user command, the user account can be assigned only to a group that already exists in the container image. Kickstart commands not listed here are allowed to be used with **ostreecontainer** command, however, they are not guaranteed to work as expected with package-based installations.

However, the following Kickstart commands are unsupported together with **ostreecontainer**:

- %packages (any necessary packages must be already available in the container image)

- url (if there is a need to fetch a **stage2** image for installation, for example, PXE installations, use **inst.stage2=** on the kernel instead of providing a url for **stage2** inside the Kickstart file)
- liveimg
- vnc
- authconfig and authselect (provide relevant configuration in the container image instead)
- module
- repo
- zipl
- zfcp

Installation of bootable OSTree native containers is not supported in interactive installations that use partial Kickstart files.

Note: When customizing a mount point, you must define the mount point in the **/mnt** directory and ensure that the mount point directory exists inside **/var/mnt** in the container image.

Jira:RHEL-2250[1]

Boot loader installation and configuration via bootupd / bootupctl in Anaconda is now available as a Technology Preview

As the **ostreecontainer** Kickstart command is now available in Anaconda as a Technology Preview, you can use it to install the operating system from an OSTree commit encapsulated in an OCI image. Anaconda automatically arranges a boot loader installation and configuration via the **bootupd**/**bootupctl** tool contained within the container image, even without an explicit boot loader configuration in Kickstart.

Jira:RHEL-17205[1]

Container-based deployments on s390x is now available as a Technology Preview

The RHEL installation program now supports deploying bootable containers in Image Mode on the **s390x** architectures by using the **ostreecontainer** Kickstart command as a Technology Preview. This enhancement removes previous limitations and ensures consistent deployment options across supported architectures. Users can now automate installations on **s390x** systems by using container-based workflows.

Jira:RHEL-63237

8.2. SECURITY

Encrypted DNS in RHEL is available as a Technology Preview

You can enable encrypted DNS to secure DNS communication that uses DNS-over-TLS (DoT). Encrypted DNS (eDNS) encrypts all DNS traffic end-to-end, with no fallback to insecure protocols, and aligns with zero trust architecture (ZTA) principles.

To perform a new installation with eDNS, specify the DoT-enabled DNS server by using the kernel command line. This ensures encrypted DNS is active during the installation process, boot time, and on

the installed system. If you require a custom CA certificate bundle, you can install it only by using the **%certificate** section in the Kickstart file. Currently, the custom CA bundle can be installed only through Kickstart installation.

On an existing system, configure NetworkManager to use a new DNS plugin, **dnsconfd**, which manages the local DNS resolver (unbound) for eDNS. Add kernel arguments to configure eDNS for the early boot process, and optionally install a custom CA bundle.

Additionally, Identity Management (IdM) deployments can also use encrypted DNS, with the integrated DNS server supporting DoT.

See Securing system DNS traffic with encrypted DNS for more details.

Jira:RHELDOCS-20059[1], Jira:RHEL-67913

New package: fips-provider-next (Technology Preview)

As a Technology Preview, this update adds a new FIPS provider that showcases future code before it obtains FIPS certification.

Jira:RHEL-96056[1]

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.

Jira:RHELPLAN-128129[1]

OpenSSL clients can use the QUIC protocol as a Technology Preview

OpenSSL can use the QUIC transport layer network protocol on the client side with the rebase to OpenSSL version 3.2.2 as a Technology Preview.

Jira:RHELDOCS-18935[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**.

Jira:RHEL-11792[1]

8.3. RHEL FOR EDGE

FDO now provides storing and querying Owner Vouchers from a SQL backend as a Technology Preview

With this Technology Preview, FDO **manufacturer-server**, **onboarding-server**, and **rendezvous-server** are available for storing and querying Owner Vouchers from a SQL backend. As a result, you can select a SQL datastore in the FDO servers options, along with credentials and other parameters, to store the Owner Vouchers.

Jira:RHELDOCS-17752[1]

8.4. SHELLS AND COMMAND-LINE TOOLS

RHEL 9.7 provides ReaR on aarch64 as a Technology Preview

RHEL 9.7 introduces the Relax and Recover (ReaR) package for the 64-bit ARM architecture (**aarch64**) as a Technology Preview. ReaR is a disaster recovery tool that produces a bootable image that you can use to restore the system from a backup. You can currently use the following output methods with ReaR on **aarch64**: ISO, USB, and PXE.

For more information about ReaR, see the article What is Relax and Recover(ReaR) and how to use it for disaster recovery?

Jira:RHEL-56045[1]

8.5. INFRASTRUCTURE SERVICES

libabigail: Flexible array conversion warning-suppression available as a Technology Preview

As a Technology Preview, when comparing binaries, you can suppress warnings related to fake flexible arrays that were converted to true flexible arrays by using the following suppression specification:

```
[suppress_type]
      type_kind = struct
      has_size_change = true
      has_strict_flexible_array_data_member_conversion = true
```

Jira:RHEL-16629[1]

8.6. NETWORKING

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.

Jira:RHEL-88552[1]

KTLS available as a Technology Preview

In RHEL, Kernel Transport Layer Security (KTLS) is provided as a Technology Preview. KTLS handles TLS records using the symmetric encryption or decryption algorithms in the kernel for the AES-GCM cipher. KTLS also includes the interface for offloading TLS record encryption to Network Interface Controllers (NICs) that provides this functionality.

Note that specific uses cases of kernel TLS offload might have a higher support status. For details see the release notes in the New features chapter.

Jira:RHEL-88551[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.

Jira:RHEL-88550

NetworkManager and the Nmstate API support MACsec hardware offload

You can use both NetworkManager and the Nmstate API to enable MACsec hardware offload if the hardware supports this feature. As a result, you can offload MACsec operations, such as encryption, from the CPU to the network interface controller.

Note that this feature is an unsupported Technology Preview.

Jira:RHEL-24337

NetworkManager enables configuring HSR and PRP interfaces

High-availability Seamless Redundancy (HSR) and Parallel Redundancy Protocol (PRP) are network protocols that provide seamless failover against failure of any single network component. Both protocols are transparent to the application layer, meaning that users do not experience any disruption in communication or any loss of data, because a switch between the main path and the redundant path happens very quickly and without awareness of the user. Now it is possible to enable and configure HSR and PRP interfaces using the **NetworkManager** service through the **nmcli** utility and the DBus message system.

Jira:RHEL-5852

UDP encapsulation in packet offload mode is now available as a Technology Preview

With IPsec packet offload, the kernel can offload the entire IPsec encapsulation process to a NIC to reduce the workload. With this update, the packet offload has been improved by supporting User Datagram Protocol (UDP) encapsulation of **ipsec** tunnels when in packet offload mode.

Jira:RHEL-30141[1]

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 Internet Protocol (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.

Jira:RHELPLAN-102815[1]

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.

Jira:RHELPLAN-129881[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_af

Marvell OcteonTX2 RVU Admin Function driver

rvu_nicpf

Marvell OcteonTX2 NIC Physical Function driver

rvu_nicvf

Marvell OcteonTX2 NIC Virtual Function driver

Jira:RHELPLAN-108169[1]

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.

Jira:RHELPLAN-154595[1]

kTLS was updated to version 6.12

The kernel Transport Layer Security (KTLS) functionality is a Technology Preview. In RHEL 9.6, we updated kTLS to the 6.12 upstream version.

Jira:RHELPLAN-153754[1]

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 redundancy with zero-time recovery in Ethernet networks.

Jira:RHELDOCS-20472[1]

8.7. KERNEL

python-drgn available as a Technology Preview

The **python-drgn** package brings an advanced debugging utility, which adds emphasis on programmability. You can use its Python command-line interface to debug both the live kernels and the kernel dumps. Additionally, **python-drgn** offers scripting capabilities for you to automate debugging tasks and conduct intricate analysis of the Linux kernel.

Jira:RHEL-6973[1]

The IAA crypto driver is now available as a Technology Preview

The Intel® In-Memory Analytics Accelerator (Intel® IAA) is a hardware accelerator that provides very high throughput compression and decompression combined with primitive analytic functions.

The **iaa_crypto** driver, which offloads compression and decompression operations from the CPU, has been introduced in RHEL 9.4 as a Technology Preview. It supports compression and decompression compatible with the DEFLATE compression standard described in RFC 1951. The **iaa_crypto** driver is designed to work as a layer underneath higher-level compression devices such as **zswap**.

For details about the IAA crypto driver, see:

- Intel® In-Memory Analytics Accelerator (Intel® IAA) User Guide
- IAA Compression Accelerator Crypto Driver

Jira:RHEL-20145[1]

The Neural Processing Unit (NPU) kernel for the RHEL Kernel is available as a Technology Preview on Intel Arrow Lake-based systems

In RHEL 9.6, the kernel introduces the Neural Processing Unit (NPU) as a Technology Preview. NPUs are special chips used for artificial intelligence (AI) and machine learning (ML) tasks on the systems. The kernel in RHEL 9.6 includes the initial driver for Intel NPUs and support infrastructure required to use the NPUs for AI/ML tasks.

Jira:RHEL-38583[1]

The Red Hat Enterprise Linux for Real Time on ARM64 is now available as a Technology Preview

With this Technology Preview, the Red Hat Enterprise Linux for Real Time is now enabled for ARM64. The ARM64 is enabled on ARM (AARCH64), for both 4k and 64k ARM kernels.

Jira:RHELDOCS-19635[1]

Boot from NVMe/TCP is available as a Technology Preview

On systems that boot from SAN over NVMe-TCP, you can use **kdump** to write crash dumps to an NVMe namespace. This update fixes failures that occurred when **kdump** attempted to dump to the NVMe namespace. As a result, panic dumps succeed on these systems, improving recovery and reducing downtime in SAN-based environments.

Jira:RHEL-33413[1]

8.8. FILE SYSTEMS AND STORAGE

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.

Jira:RHELPLAN-102321[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**).

Jira:RHELPLAN-58357[1]

NVMe/TCP using TLS is available as a Technology Preview

Encrypting Non-volatile Memory Express (NVMe) over TCP (NVMe/TCP) network traffic using TLS configured with Pre-Shared Keys (PSK) has been added as a Technology Preview in RHEL 9.6. For instructions, see Configuring an NVMe/TCP host using TLS with Pre-Shared-Keys .

Jira:RHEL-9301[1]

xfs_scrub utility is available as a Technology Preview

You can check all the metadata on a mounted XFS file system by using the **xfs_scrub** utility as a Technology Preview. It functions similarly to the **xfs_repair -n** command for an unmounted XFS filesystem. For details, see the **xfs_scrub(8)** man page on your system. Note that currently only the

scrub feature is available in RHEL 10 kernels and online repair is not enabled.

Jira:RHELDOCS-21350[1]

8.9. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

A new nodejs:22 module stream available as a Technology Preview

A new module stream, **nodejs:22**, is now available as a Technology Preview. A future update will provide a Long Term Support (LTS) version of **Node.js 22**, which will be fully supported.

Node.js 22 included in RHEL 9.5 provides numerous new features, bug fixes, security fixes, and performance improvements over **Node.js 20** available since RHEL 9.3.

Notable changes include:

- The **V8** JavaScript engine has been upgraded to version 12.4.
- The **V8 Maglev** compiler is now enabled by default on architectures where it is available (AMD and Intel 64-bit architectures and the 64-bit ARM architecture).
- **Maglev** improves performance for short-lived CLI programs.
- The **npm** package manager has been upgraded to version 10.8.1.
- The **node --watch** mode is now considered stable. In **watch** mode, changes in watched files cause the **Node.js** process to restart.
- The browser-compatible implementation of **WebSocket** is now considered stable and enabled by default. As a result, a WebSocket client to Node.js is available without external dependencies.
- **Node.js** now includes an experimental feature for execution of scripts from **package.json**. To use this feature, run the **node --run <script-in-package.json>** command.

To install the **nodejs:22** module stream, enter:

```
# dnf module install nodejs:22
```

If you want to upgrade from the **nodejs20** 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.

Jira:RHEL-35990

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.

Jira:RHELPLAN-88788[1]

A new nodejs:24 module stream is available as a Technology Preview

A new **nodejs:24** module stream is available as a Technology Preview in Red Hat Enterprise Linux 9.7. This update introduces Node.js 24, which provides new features, bug fixes, security updates, and performance improvements compared to Node.js 22 included in RHEL 9.6.

To install the **nodejs:24** module, enter:

```
# dnf module install nodejs:24
```

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

Jira:RHEL-90821

8.10. COMPILERS AND DEVELOPMENT TOOLS

eu-stacktrace available as a Technology Preview

The **eu-stacktrace** utility, which has been distributed through the **elfutils** package since version 0.192, is available as a Technology Preview feature. **eu-stacktrace** is a prototype utility that uses the **elfutils** toolkit's unwinding libraries to support a sampling profiler to unwind frame pointer-less stack sample data.

Jira:RHELDOCS-19072[1]

8.11. IDENTITY MANAGEMENT

DNS over TLS (DoT) in IdM deployments is available as a Technology Preview

Encrypted DNS using DNS over TLS (DoT) is now available as a Technology Preview in Identity Management (IdM) deployments. You can now encrypt all DNS queries and responses between DNS clients and IdM DNS servers.

To start using this functionality, install the **ipa-server-encrypted-dns** package for IdM servers and replicas, and the **ipa-client-encrypted-dns** package for IdM clients. Administrators can enable DoT during the installation using the **--dns-over-tls** option.

IdM configures Unbound as a local caching resolver and BIND to receive DoT requests. This functionality is available through the command-line interface (CLI) and non-interactive installations of IdM.

To configure DoT, new options were added to installation utilities for IdM servers, replicas, clients, and the integrated DNS service:

- **--dot-forwarder** to specify an upstream DoT-enabled DNS server.

- **--dns-over-tls-key** and **--dns-over-tls-cert** to configure DoT certificates.
- **--dns-policy** to set a DNS security policy to either allow fallback to unencrypted DNS or enforce strict DoT usage.

By default, IdM uses **relaxed** DNS policy, which allows fallback to unencrypted DNS. You can enforce encrypted-only communication using the new **--dns-policy** option with the **enforced** setting.

You can also enable DoT on an existing IdM deployment by reconfiguring the integrated DNS service using **ipa-dns-install** with the new DoT options.

See Securing DNS with DoT in IdM for more details.

Jira:RHEL-67913[1], Jira:RHELDOCS-20059

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.

Jira:RHELPLAN-121751[1]

Encrypted DNS with DoT is now available in ansible-freeipa installations of IdM as a Technology Preview

You can now use Ansible to ensure that all DNS queries and responses between DNS clients and Identity Management (IdM) DNS servers are encrypted. Encrypted DNS using DNS over TLS (DoT) has been available as a Technology Preview in IdM deployments since RHEL 10. In RHEL 10.1, the functionality is available as a Technology Preview in the **freeipa.ansible_freeipa** collection.

To enable DoT during a deployment of IdM by using **ansible-freeipa** use the following options:

- **ipaserver_dns_over_tls** with the **freeipa.ansible_freeipa.ipaserver** role for a new server.
- **ipareplica_dns_over_tls** with the **freeipa.ansible_freeipa.ipareplica** role for a replica.
- **dot_forwarder** to specify an upstream DoT-enabled DNS server.
- **dns_over_tls_key** and **dns_over_tls_cert** to configure DoT certificates.

Additionally, you can set the **dns_policy** variable to enforce DoT-only communication, overriding the default behavior that allows fallback to unencrypted DNS.

Jira:RHELDOCS-20258[1]

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.

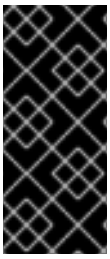

IMPORTANT

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
```

Jira:RHELPLAN-121754[1]

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

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

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

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

Jira:RHELPLAN-27737[1]

8.13. THE WEB CONSOLE

The RHEL web console can now manage WireGuard connections

Starting with RHEL 9.4, you can use the RHEL web console to create and manage WireGuard VPN connections. Note that, both the WireGuard technology and its web console integration are unsupported Technology Previews.

Jira:RHELDOCS-17520[1]

8.14. 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:RHELDOCS-17040[1]

AMD SEV, SEV-ES, and SEV-SNP 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.

RHEL 9.5 and later also provides the Secure Nested Paging (SEV-SNP) feature as Technology Preview. SNP enhances SEV and SEV-ES by improving its memory integrity protection, which helps prevent hypervisor-based attacks, such as data replay or memory re-mapping.

Note that SEV and SEV-ES work only on the 2nd generation of AMD EPYC CPUs (codenamed Rome) or later. Similarly, SEV-SNP works only on 4rd generation AMD EPYC CPUs (codenamed Genoa) or later. Also note that RHEL 9 includes SEV, SEV-ES, and SEV-SNP encryption, but not the SEV, SEV-ES, and SEV-SNP security attestation and live migration.

Jira:RHELPLAN-65217[1]

CPU clusters on 64-bit ARM

As a Technology Preview, you can now create KVM virtual machines that use multiple 64-bit ARM CPU clusters in their CPU topology.

Jira:RHEL-7043[1]

New package: trustee-guest-components

As a Technology Preview, this update adds the **trustee-guest-components** package. This makes it possible for confidential virtual machines to attest themselves and get confidential resources from a Trustee server.

Jira:RHEL-68141[1]

8.15. CONTAINERS

The podman-machine command is unsupported

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

Jira:RHELDOCS-16861[1]

A new rhel9/rhel-bootc container image is available as a Technology Preview

The **rhel9/rhel-bootc** container image is now available in the Red Hat Container Registry as a Technology Preview. With the RHEL bootable container images, you can build, test, and deploy an

operating system exactly as a container. The RHEL bootable container images differ from the existing application Universal Base Images (UBI) thanks to the following enhancements: RHEL bootable container images contain additional components necessary to boot, such as, kernel, initrd, bootloader, firmware, between others. There are no changes to existing container images. For more information, see Red Hat Ecosystem Catalog .

Jira:RHELDOCS-17803[1]

Partial pulls for zstd:chunked are available as a Technology Preview

You can pull only the changed parts of the container images compressed with the **zstd:chunked** format, reducing network traffic and necessary storage. You can enable partial pulls by adding the **enable_partial_images = "true"** setting to the **/etc/containers/storage.conf** file. This functionality is available as a Technology Preview.

Jira:RHEL-32267

The podman artifact command is available as a Technology Preview

The **podman artifact** command, which you can use to work with OCI artifacts at the command-line level, is available as a Technology Preview. For further information, reference the man page.

Jira:RHEL-70217

Podman compatibility with Docker API is available as a Technology Preview

Podman supports the following Docker API versions as a Technology Preview:

- Docker API 1.41
- Docker API 1.43

Jira:RHEL-88121

CHAPTER 9. DEPRECATED FUNCTIONALITIES

Deprecated devices are fully supported, which means that they are tested and maintained, and their support status remains unchanged within Red Hat Enterprise Linux 9. However, these devices will likely not be supported in the next major version release, and are not recommended for new deployments on the current or future major versions of RHEL.

For the most recent list of deprecated functionality within a particular major release, see the latest version of release documentation. For information about the length of support, see Red Hat Enterprise Linux Life Cycle and Red Hat Enterprise Linux Application Streams Life Cycle .

A package can be deprecated and not recommended for further use. Under certain circumstances, a package can be removed from the 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 .

9.1. 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**
- **nvdimm**

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.

Jira:RHELPLAN-60153[1]

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:RHELDOCS-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.

Jira:RHELPLAN-168262[1]

Anaconda built-in help has been deprecated

The built-in documentation from spokes and hubs of all Anaconda user interfaces, which is available during Anaconda installation, has been deprecated. As a replacement, the Anaconda user interfaces will be self-descriptive and users can refer to the official RHEL documentation in future major RHEL releases.

Jira:RHELDOCS-17309[1]

Support for NVDIMM devices has been deprecated

Previously, the installation program allowed reconfiguring NVDIMM devices during installation. This support for NVDIMM devices during the Kickstart and GUI installation has been deprecated, and will be removed in the next major RHEL release. The NVDIMM devices in the sector mode will still be visible and usable in the installation program.

Jira:RHELDOCS-17702

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

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

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

9.2. SECURITY

X25519-MLKEM768 deprecated and aliased to MLKEM768-X25519 in crypto-policies

The **X25519-MLKEM768** value in system-wide cryptographic policies is deprecated and aliased to the **MLKEM768-X25519** value. This unifies the concatenation order, allowing both variants to work.

Jira:RHEL-103793

Keylime policy management scripts are deprecated and replaced with keylime-policy

In RHEL 9.6, Keylime is provided with the **keylime-policy** tool, which replaces the following policy management scripts:

- **keylime_convert_runtime_policy**
- **keylime_create_policy**
- **keylime_sign_runtime_policy**
- **create_mb_refstate**
- **create_allowlist.sh**

These scripts have been deprecated and will be removed in a future major version of RHEL.

Jira:RHELDOCS-19815[1]

OVAL deprecated in vulnerability scanning applications

The Open Vulnerability Assessment Language (OVAL) data format, which provides declarative security data processed by the OpenSCAP suite, is deprecated and will be removed in a future major release. Red Hat continues to provide declarative security data in the Common Security Advisory Framework (CSAF) format, which is the successor of OVAL.

For more information, see the OVAL v2 Announcement.

Alternatively, you can us Red Hat Lightspeed for RHEL vulnerability service, for more information, follow Assessing and Monitoring Security Vulnerabilities on RHEL Systems .

Jira:RHELDOCS-17532[1]

libgcrypt is deprecated

The Libgcrypt cryptographic library provided by the **libgcrypt** package is deprecated and may be removed in a future major release. Instead, use the libraries listed in the RHEL core cryptographic components article (Red Hat Knowledgebase).

Jira:RHELDOCS-17508[1]

fips-mode-setup is deprecated

The **fips-mode-setup** tool, which switches the system to FIPS mode, is deprecated in RHEL 9. You can still use the **fips-mode-setup** command to check whether FIPS mode is enabled.

To operate a system compliant with FIPS 140, install a system in FIPS mode in one of the following ways:

- Add the **fips=1** option to the kernel command line during the RHEL installation. See the Customizing boot options chapter in the Interactively installing RHEL from installation media document for more information.
- Create a FIPS-enabled image with RHEL image builder by adding the **fips=yes** directive to the **[customizations]** section of its blueprint.
- Create a disk image with the **bootc-image-builder** tool or install the system by using the **bootc install-to-disk** tool with a Containerfile that follows the example in the *Using image mode for RHEL* document to add the **fips=1** kernel command line flag and switch the system-wide cryptographic policy to **FIPS**.

The **fips-mode-setup** tool will be removed in the next major release.

Jira:RHELDOCS-19284

Using update-ca-trust without arguments is deprecated

Previously, the command **update-ca-trust** updated the system certificate authority (CA) store regardless of the arguments entered. This update introduces the **extract** subcommand for updating the CA store. You can also specify the location to which the CA certificates are extracted by using the **--output** argument. For compatibility with earlier versions of RHEL, entering **update-ca-trust** to update the CA store with any argument other than **-o** or **--help**, and even without any argument, is still supported for the duration of RHEL 9, but will be removed by the next major release. Update your calls to **update-ca-trust extract**.

Jira:RHEL-54695[1]

CAfile pointing to trusted root certificate files in Stunnel clients is deprecated

If Stunnel is configured in client mode, the **CAfile** directive can point to a file that contains trusted root certificates in the **BEGIN TRUSTED CERTIFICATE** format. This method is deprecated and might be removed in a future major version. In a future version, **stunnel** will pass the value of the **CAfile** directive to a function that does not support the **BEGIN TRUSTED CERTIFICATE** format. As a consequence, if you use **CAfile = /etc/pki/tls/certs/ca-bundle.trust.crt**, change the location to **CAfile = /etc/pki/tls/certs/ca-bundle.crt**.

Jira:RHEL-52317[1]

DSA and SEED algorithms have been deprecated in NSS

The Digital Signature Algorithm (DSA), which was created by the National Institute of Standards and Technology (NIST) and is now completely deprecated by NIST, is deprecated in the Network Security Services (NSS) cryptographic library. You can instead use algorithms such as RSA, ECDSA, and EdDSA.

The SEED algorithm, which was created by the Korea Information Security Agency (KISA) and has been previously disabled upstream, is deprecated in the NSS cryptographic library.

Jira:RHELDOCS-19004[1]

pam_ssh_agent_auth is deprecated

The **pam_ssh_agent_auth** package is deprecated and might be removed in a future major release.

Jira:RHELDOCS-18312[1]

compat-openssl11 is deprecated

The compatibility library for OpenSSL 1.1, **compat-openssl11**, is now deprecated, and it might be removed in a future major release. OpenSSL 1.1 is no longer maintained upstream and applications that use the OpenSSL TLS toolkit should be migrated to version 3.x.

Jira:RHELDOCS-18480[1]

SHA-1 is deprecated at SECLEVEL=2 in OpenSSL

The use of the SHA-1 algorithm at **SECLEVEL=2** is deprecated in OpenSSL and might be removed in a future major release.

Jira:RHELDOCS-18701[1]

OpenSSL Engines API is deprecated in Stunnel

The use of the OpenSSL Engines API in Stunnel is deprecated and will be removed in a future major release. The most common use is to access hardware security tokens that use PKCS#11 through the **openssl-pkcs11** package. As a replacement, you can use **pkcs11-provider**, which uses the new OpenSSL Providers API.

Jira:RHELDOCS-18702[1]

OpenSSL Engines are deprecated

OpenSSL Engines are deprecated and will be removed in the near future. Instead of using engines, you can use the **pkcs11-provider** as a replacement.

Jira:RHELDOCS-18703[1]

DSA is deprecated in GnuTLS

The Digital Signature Algorithm (DSA) is deprecated in the GnuTLS secure communications library and will be removed in a future major version of RHEL. DSA was previously deprecated by the National Institute of Standards and Technology (NIST), and is not considered secure. You can use ECDSA instead to ensure compatibility with future versions.

Jira:RHELDOCS-19224[1]

scap-workbench is deprecated

The **scap-workbench** package is deprecated. The **scap-workbench** graphical utility was designed to perform configuration and vulnerability scans on a single local or remote system. As an alternative, you can scan local systems for configuration compliance by using the **oscap** command and remote systems by using the **oscap-ssh** command. For more information, see Configuration compliance scanning.

Jira:RHELDOCS-19028[1]

oscap-anaconda-addon is deprecated

The **oscap-anaconda-addon**, which provided means to deploy baseline-compliant RHEL systems by using the graphical installation, is deprecated. As an alternative, you can build RHEL images that comply with a specific standard by Creating pre-hardened images with RHEL image builder OpenSCAP integration.

Jira:RHELDOCS-19029[1]

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.

Jira:RHELPLAN-112355[1]

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.

Jira:RHELPLAN-148207[1]

OpenSSL deprecates the Engines API

The OpenSSL 3.0 TLS toolkit deprecated the Engines API. The Engines interface is superseded by the Providers API. The migration of applications and existing engines to Providers is underway. The deprecated Engines API may be removed in a future major release.

Jira:RHELDOCS-17958[1]

openssl-pkcs11 is now deprecated

As a part of the ongoing migration of deprecated OpenSSL engines to the Providers API, the **pkcs11-provider** package replaces the **openssl-pkcs11** package (**engine_pkcs11**). The **openssl-pkcs11** package is now deprecated. The **openssl-pkcs11** package may be removed in a future major release.

Jira:RHELDOCS-16716[1]

RHEL 8 and 9 OpenSSL certificate and signing containers are now deprecated

The OpenSSL portable certificate and signing containers available in the **ubi8/openssl** and **ubi9/openssl** repositories in the Red Hat Ecosystem Catalog are now deprecated due to low demand.

Jira:RHELDOCS-17974[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.

Jira:RHELPLAN-94096[1]

/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 displaying the **/proc/sys/crypto/fips_enabled** file.

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.

Jira:RHELPLAN-106338[1]

OpenSSL does not accept explicit curve parameters in FIPS mode

Elliptic curve cryptography parameters, private keys, public keys, and certificates that specified explicit curve parameters no longer work in FIPS mode. Specifying the curve parameters using ASN.1 object identifiers, which use one of the FIPS-approved curves, still works in FIPS mode.

Jira:RHELPLAN-113856[1]

OpenSSL rejects RSA signatures with X9.31 padding in FIPS mode

Because X9.31 RSA signatures were removed from the FIPS 186-5 standard, OpenSSL no longer supports signing or signature verification with RSA keys with X9.31 padding in FIPS mode.

Jira:RHELPLAN-139207[1]

9.3. RHEL FOR EDGE

Ignition has been deprecated for image mode for RHEL for Edge images

The Ignition tool, used to inject the user configuration into the Simplified Installer, AMI, and VMDK RHEL for Edge images types at an early stage of the boot process, has been deprecated in RHEL 9 and might be removed in a future major release.

Jira:RHELDOCS-19754[1]

9.4. SUBSCRIPTION MANAGEMENT

Several subscription-manager modules have been deprecated

Because of a simplified customer experience in Red Hat subscription services, which have transitioned to the Red Hat Hybrid Cloud Console and to account level subscription management with Simple Content Access, the following modules have been deprecated and will be removed in a future major release:

- **addons**
- **attach**
- **auto-attach**
- **import**
- **remove**
- **redeem**
- **role**
- **service-level**
- **syspurpose addons**
- **usage** For more information about these transitions, see the Transition of Red Hat's subscription services to the Red Hat Hybrid Cloud Console article.

Jira:RHEL-29178

9.5. SOFTWARE MANAGEMENT

The numberless %patch syntax has been deprecated

Using the **%patch** directive without a number specified as a shorthand for **%patch 0** to apply the **zero-th** patch has been deprecated. Therefore, if you want to use **%patch**, a warning message suggests you to use the explicit syntax, for example, **%patch 0** or **%patch -P 0** to apply the **zero-th** patch.

Jira:RHELDOCS-19810[1]

The DNF debug plug-in has been deprecated

The DNF **debug** plug-in, which includes the **dnf debug-dump** and **dnf debug-restore** commands, has been deprecated and will be removed from the **dnf-plugins-core** package in the next major RHEL release.

Jira:RHELDOCS-18592[1]

The support for libreport has been deprecated

The support for the **libreport** library has been deprecated and will be removed from DNF in the next major RHEL release.

Jira:RHELDOCS-18593[1]

9.6. SHELLS AND COMMAND-LINE TOOLS

The dump utility from the dump package has been deprecated

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

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

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

Jira:RHELPLAN-94704[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

The %vmeff metric from the sysstat package has been deprecated

The **%vmeff** metric from the **sysstat** package to measure the page reclaim efficiency will no longer be supported in a future major version of RHEL. 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.

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:RHELDOCS-17015[1]

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=...**, 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:RHELDOCS-18049[1]

cgroupsv1 is now deprecated in RHEL 9

The **cgroups** is a kernel subsystem used for process tracking, system resource allocation and partitioning. Systemd service manager supports booting in the cgroups **v1** mode as well as in cgroups **v2** mode. In Red Hat Enterprise Linux 9, the default mode is **v2**. In Red Hat Enterprise Linux 10, systemd will not support booting in the cgroups **v1** mode and only cgroups **v2** mode will be available.

Jira:RHELDOCS-17545[1]

9.7. INFRASTRUCTURE SERVICES

Various packages are now deprecated in infrastructure services

The following packages are deprecated in RHEL 9 and will not be distributed in later major versions of RHEL:

- **sendmail**
- **libotr**
- **mod_security**
- **spamassassin**
- **redis**
- **dhcp**
- **xsane**

Jira:RHEL-22385[1]

9.8. NETWORKING

ipset has been deprecated

In RHEL 9, the **ipset** utility is deprecated and is planned to be removed in a future major release. Red Hat will provide bug fixes and support for this feature during the current release lifecycle, but this feature will no longer receive enhancements. As an alternative to **ipset**, you can use the **nftables** sets functionality instead.

Jira:RHELDOCS-20146[1]

The Soft-iWARP driver is deprecated

RHEL 9 provides the Soft-iWARP driver as an unsupported Technology Preview. Starting with RHEL 9.5, this driver is deprecated and will be removed in RHEL 10.

Jira:RHELDOCS-18699[1]

The dhcp-client package is deprecated

Previously, you could configure NetworkManager in RHEL 9 to use a DHCP client from the **dhcp-client** package. However, the option to use the **dhclient** utility is now deprecated and results in a warning being displayed at the NetworkManager startup. To configure NetworkManager as described above, switch to the internal DHCP library. In RHEL 10, the **dhcp-client** package is no longer available and the applications configured to use the **dhclient** utility use the internal DHCP library instead.

Jira:RHEL-24622

The perl(Mail::Sender) module is now deprecated

The **perl(Mail::Sender)** module is now deprecated and will be removed from the next major release without any replacement. As a result, the **checkbandwidth** script from **net-snmp-perl** package does not support email alerts when bandwidth high/low levels for a host or interface are reached.

Jira:RHELDOCS-18959[1]

libdb has been deprecated

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 following Red Hat Knowledgebase articles:

- How to migrate from libdb to a different key-value database
- Available replacements for the deprecated Berkeley DB (libdb) in RHEL

Jira:RHELPLAN-67314[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.

Jira:RHELPLAN-69554[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 .

Jira:RHELPLAN-58745[1]

The iptables back end in firewalld is deprecated

In RHEL 9, the **iptables** framework is deprecated. As a consequence, the **iptables** back end 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.

Jira:RHELPLAN-122745[1]

The firewalld lockdown feature is deprecated.

The lockdown feature in **firewalld** is deprecated because it cannot prevent processes that are running as **root** from adding themselves to the allow list. The lockdown feature may be removed in a future major RHEL release.

Jira:RHEL-17708

The connection.master, connection.slave-type, and connection.autoconnect-slaves properties are deprecated

Red Hat is committed to using conscious language. Therefore, the **connection.master**, **connection.slave-type**, and **connection.autoconnect-slaves** properties were renamed. To ensure backward compatibility, aliases have been created that map the old property names to the new ones:

- **connection.master** is an alias for **connection.controller**
- **connection.slave-type** is an alias for **connection.port-type**
- **connection.autoconnect-slaves** is an alias for **connection.autoconnect-ports**

Note that the **connection.master**, **connection.slave-type**, and **connection.autoconnect-slaves** aliases are deprecated and will be removed in a future RHEL version.

Jira:RHEL-17619[1]

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]

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 .

Jira:RHELPLAN-113659[1]

Client-side and server-side DHCP packages are deprecated

Internet Systems Consortium (ISC) has announced the end of maintenance for ISC DHCP as of the end of 2022. As a result, Red Hat has decided to deprecate the use of client-side and server-side DHCP packages in RHEL 9 and not to distribute them in later major versions of RHEL. Customers must prepare for the transition to available alternatives, such as **dhcpcd** and **ISC Kea**.

Jira:RHELDOCS-17135[1]

9.9. KERNEL

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? .

Jira:RHELPLAN-129876[1]

The deprecated --token option of subscription-manager register will stop working at the end of November 2024

The deprecated **--token=<TOKEN>** option of the **subscription-manager register** command will no longer be a supported authentication method from the end of November 2024. The default entitlement server, **subscription.rhsm.redhat.com**, will no longer be allowing token-based authentication. As a consequence, if you use **subscription-manager register --token=<TOKEN>**, the registration will fail with the following error message:

> Token authentication not supported by the entitlement server

To register your system, use other supported authorization methods, such as including paired options **--username** / **--password** OR **--org** / **--activationkey** with the **subscription-manager register** command.

Jira:RHELPLAN-146101[1]

9.10. FILE SYSTEMS AND STORAGE

Support for the block translation table driver has been deprecated

Support for the block translation table driver (btt.ko) has been deprecated and will be removed in the future major RHEL release. Red Hat will provide bug fixes and support for configuring Non-Volatile Dual In-line Memory Modules (NVDIMM) namespaces by using sector mode during the current release lifecycle. However, this feature will no longer receive enhancements and will be removed.

Jira:RHELDOCS-19716[1]

The nvme_core.multipath parameter is deprecated

In RHEL 9.6, the **nvme_core.multipath** parameter is deprecated and is planned to be removed in a future release. Red Hat will provide bug fixes and support for this feature during the current release lifecycle, but this feature will no longer receive enhancements and will be removed in a future major release.

Jira:RHELDOCS-19809[1]

Support for NVMe devices has been deprecated from the lsscsi package

Support for Non-volatile Memory Express (NVMe) devices has been deprecated and will be removed from the **lsscsi** package in the future major RHEL release. Use native tools such as **nvme-cli**, **lsblk**, and **blkid** instead.

Jira:RHELDOCS-19068[1]

Support for NVMe devices has been deprecated from the sg3_utils package

Support for Non-volatile Memory Express (NVMe) devices has been deprecated and will be removed from the **sg3_utils** package in the future major RHEL release. You can use native tools (**nvme-cli**) instead.

Jira:RHELDOCS-19069[1]

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.

Jira:RHELPLAN-107107[1]

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:RHELDOCS-16432[1]

The md-linear, md-faulty, and md-multipath modules have been deprecated

The following MD RAID kernel modules have been deprecated and will be removed in a future major RHEL release:

- **CONFIG_MD_LINEAR** or **md-linear** to concatenate multiple drives so that when a single member disk becomes full, data are written to the next disk until all disks are full.
- **CONFIG_MD_FAULTY** or **md-faulty** to test a block device that occasionally returns read or write errors.
- **CONFIG_MD_MULTIPATH** or **md-multipath** to take advantage of hardware supporting more than one I/O path to individual LUNs (disk drives). **md-multipath** allows the data availability in case of a hardware failure or individual path saturation.

Jira:RHEL-30730[1]

The VDO sysfs parameters have been deprecated

The Virtual Data Optimizer (VDO) **sysfs** parameters have been deprecated and will be removed in a future major RHEL release. Except for **log_level**, all module-level **sysfs** parameters for the **kvdo** module will be removed. For individual **dm-vdo** targets, all **sysfs** parameters specific to VDO will also be removed. There is no change for the parameters that are common to all DM targets. Configuration values for **dm-vdo** targets, which are currently set by updating the removed module-level parameters, can no longer be changed.

Statistics and configuration values for **dm-vdo** targets will no longer be accessible through **sysfs**. But these values are still accessible by using **dmsetup message stats**, **dmsetup status**, and **dmsetup table** dmsetup commands

Jira:RHEL-30525

9.11. HIGH AVAILABILITY AND CLUSTERS

Deprecated high availability features

The following features were deprecated as of Red Hat Enterprise Linux 9.5 and will be removed in the next major release. The **pcs** command-line interface produces a warning when you attempt to configure a system with these features.

- Configuring a **score** parameter in order constraints
- Use of the **rkt** container engine in bundles
- Support for **upstart** and **nagios** resources
- The **monthdays**, **weekdays**, **weekyears**, **yearsdays** and **moon** date specification options for configuring Pacemaker rules
- The **yearsdays** and **moon** duration options for configuring Pacemaker rules

Jira:RHEL-34781

Resilient Storage Add-On has been deprecated

The Red Hat Enterprise Linux (RHEL) Resilient Storage Add-On has been deprecated as of RHEL 9. The Resilient Storage Add-On will no longer be supported starting with Red Hat Enterprise Linux 10 and any subsequent releases after RHEL 10. The RHEL Resilient Storage Add-On will continue to be supported with earlier versions of RHEL (7, 8, 9) and throughout their respective maintenance support lifecycles.

Jira:RHELDOCS-19022[1]

9.12. COMPILERS AND DEVELOPMENT TOOLS

Redis will be replaced with Valkey in Grafana, PCP, and grafana-pcp

The **Redis** key-value store has been deprecated and will be replaced with **Valkey** in the next major version of RHEL. As a result, **Grafana**, PCP, and the **grafana-pcp** plug-in will use **Valkey** to store data instead of **Redis** in RHEL 10.

Jira:RHELDOCS-18207[1]

HTML content of llvm-doc is deprecated

The HTML content of the **llvm-doc** package will be removed in a future RHEL release and replaced with a single HTML file pointing to online documentation at llvm.org. Users of **llvm-doc** that do not have network access will need an alternative way to access LLVM documentation.

Jira:RHELDOCS-19013[1]

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.

Jira:RHELPLAN-129104[1]

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.

Jira:RHELPLAN-123778[1]

32-bit packages are deprecated

Linking against 32-bit multilib packages is deprecated. The ***.i686** packages will remain supported for the life cycle of Red Hat Enterprise Linux 9, but will be removed in the next major version of RHEL.

Jira:RHELDOCS-17917[1]

9.13. IDENTITY MANAGEMENT

The dnssec-enable: no; option has been deprecated

The **dnssec-enable: no;** option in the **/etc/named/ipa-options-ext.conf** file has been deprecated and will be removed in a future major version of RHEL. DNS Security Extensions (DNSSEC) are enabled by default and disabling them will not be possible. The **dnssec-validation: no;** option still continues to be available.

Jira:RHELDOCS-20464[1]

nsslapd-subtree-rename-switch is deprecated in 389-ds-base

Before this update, you could configure Directory Server to prevent moving entries between sub-trees in a database. Because of the stability issues, this feature is deprecated and will be removed in a future major RHEL release.

Do not use the **nsslapd-subtree-rename-switch** parameter to deactivate moving entries between sub-trees. As an alternative, you can deactivate moving the entries by creating an access control instruction (ACI).

Jira:RHELDOCS-20337[1]

The pam_console module is deprecated

In RHEL 9.5, the **pam_console** module is deprecated and is planned to be removed in a future release. The **pam_console** module grants file permissions and authentication capabilities to users logged in at the physical console or terminals, and adjusts these privileges based on console login status and user presence. As an alternative to **pam_console**, you can use the **systemd-logind** system service instead. For configuration details, see the **logind.conf(5)** man page.

Jira:RHELDOCS-18158[1]

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.

Jira:RHELPLAN-88246[1]

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
```

3. To restore caching and synchronization of user information, enable the integration between **shadow-utils** and **sssd_cache** by creating a symbolic link:

```
# ln -s /usr/sbin/sss_cache /usr/sbin/sss_cache_shadow_utils
```

Jira:RHELPLAN-100639[1], Jira:RHEL-56352

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]

9.14. SSSD

The ad_allow_remote_domain_local_groups option has been deprecated

The **ad_allow_remote_domain_local_groups** option in **sssd.conf** has been deprecated in Red Hat Enterprise Linux (RHEL) 9.6. The **ad_allow_remote_domain_local_groups** option might be removed from a future release of RHEL.

Jira:RHELDOCS-19455[1]

The sss_ssh_knownhostsproxy tool has been deprecated

The **sss_ssh_knownhostsproxy** has been deprecated and will be replaced by a more efficient tool in RHEL 10. **sss_ssh_knownhostsproxy** will be kept for backwards compatibility in RHEL 9 and will be removed in RHEL 10. Support for the ssh **KnownHostsCommand** option will be added in a future release.

Jira:RHELDOCS-19115[1]

The enumeration feature has been deprecated for AD and IdM

The **enumeration** feature enables you to list all users or groups by using **getent passwd** or **getent group** commands without arguments for Active Directory (AD), Identity Management (IdM), and LDAP providers. Support for the **enumeration** feature has been deprecated for AD and IdM in Red Hat Enterprise Linux (RHEL) 9. The **enumeration** feature will be removed for AD and IdM in RHEL 10.

Jira:SSSD-6596

The libsss_simpleifp subpackage has been deprecated

The **libsss_simpleifp** subpackage that provides the **libsss_simpleifp.so** library has been deprecated in Red Hat Enterprise Linux (RHEL) 9. The **libsss_simpleifp** subpackage might be removed from a future release of RHEL.

Jira:SSSD-6601

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]

9.15. DESKTOP

Firefox and Thunderbird Flatpak images have been deprecated

The **rhel9/firefox-flatpak** and **rhel9/thunderbird-flatpak** Flatpak images, which are available in RHEL 9 as Technology Previews, have been deprecated and will be replaced by their RHEL 10 versions.

Jira:RHEL-91106[1]

Evince has been deprecated

Evince, a document viewer for the GNOME desktop, has been deprecated and will be removed in a future major release.

Jira:RHELDOCS-19889[1]

power-profile-daemon is deprecated

The **power-profile-daemon** package has been deprecated and is replaced by the **tuned-ppd** package. In new installations of RHEL 9.6, the **tuned-ppd** package is installed by default.

For systems updated to RHEL 9.6 from earlier versions, **power-profile-daemon** remains installed. If your scenario requires the use of **tuned-ppd** on an updated RHEL 9.6 version, install it manually:

```
# dnf install tuned-ppd
```

To verify that the package is installed, enter the following command:

```
# rpm -q tuned-ppd
tuned-ppd-2.25.1-1.el9.noarch
```

Jira:RHEL-68152

Totem media player has been deprecated

The Totem media player has been deprecated in RHEL 9.5 and will be removed in a future major release.

Jira:RHELDOCS-19050[1]

power-profiles-daemon has been deprecated

The **power-profiles-daemon** package that provides the power mode configuration in GNOME has been deprecated and will be removed in a future major release.

You can use Tuned as a replacement for power mode configuration in GNOME. You can use the **tuned-ppd** API translation daemon as a drop-in replacement for **power-profiles-dameon**.

Jira:RHELDOCS-19093[1]

gedit is deprecated

gedit, the default graphical text editor in Red Hat Enterprise Linux, has been deprecated and will be removed in a future major release. Instead, use GNOME Text Editor.

Jira:RHELDOCS-19149[1]

Qt 5 libraries have been deprecated

Qt 5 libraries have been deprecated and will be removed in a future major release. Qt 5 libraries are replaced with Qt 6 libraries, with new functionality and better support.

For more information, see Porting to Qt 6.

Jira:RHELDOCS-19133[1]

WebKitGTK has been deprecated

The WebKitGTK web browser engine has been deprecated and will be removed in a future major release.

As a consequence, you will no longer be able to build applications that depend on WebKitGTK. Desktop applications other than Mozilla Firefox can no longer display web content. There is no alternative web browser engine provided in RHEL 10.

Jira:RHELDOCS-19171[1]

Evolution has been deprecated

Evolution is a GNOME application that provides integrated email, calendar, contact management, and communications functionality. The application and its plugins has been deprecated and will be removed in a future major version. You can find an alternative in a third party source, for example on Flathub.

Jira:RHELDOCS-19147[1]

Festival has been deprecated

The Festival speech synthesizer has been deprecated and will be removed in a future major version.

As an alternative, you can use the Espeak NG speech synthesizer.

Jira:RHELDOCS-19139[1]

The Eye of GNOME has been deprecated

The Eye of GNOME (**eog**) image viewer application has been deprecated in RHEL 9.

As an alternative, you can use the Loupe application.

Jira:RHELDOCS-19135[1]

Cheese has been deprecated

The Cheese camera application has been deprecated and will be removed in a future major version.

As an alternative, you can use the Snapshot application.

Jira:RHELDOCS-19137[1]

Devhelp has been deprecated

Devhelp, a graphical developer tool for browsing and searching API documentation, has been deprecated and will be removed in a future major version. You can now find API documentation online in specific upstream projects.

Jira:RHELDOCS-19154[1]

gtkmm based on GTK 3 has been deprecated

gtkmm is a C++ interface for the GTK graphical toolkit. The **gtkmm** version that was based on GTK 3 has been deprecated with all its dependencies and will be removed in a future major version. To access **gtkmm** in RHEL 10, migrate to the **gtkmm** version based on GTK 4.

Jira:RHELDOCS-19143[1]

Inkscape has been deprecated

The Inkscape vector graphics editor has been deprecated and will be removed in a future major version.

Jira:RHELDOCS-19151[1]

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]

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]

TigerVNC is deprecated

The TigerVNC remote desktop solution is now deprecated. It will be removed in a future major RHEL release and replaced by a different remote desktop solution.

TigerVNC provides the server and client implementation of the Virtual Network Computing (VNC) protocol in RHEL 9.

The following packages are deprecated:

- **tigervnc**

- **tigervnc-icons**
- **tigervnc-license**
- **tigervnc-selinux**
- **tigervnc-server**
- **tigervnc-server-minimal**
- **tigervnc-server-module**

The **Connections** application (**gnome-connections**) continues to be supported as an alternative VNC client, but it does not provide a VNC server.

Jira:RHELDOCS-17782[1]

9.16. GRAPHICS INFRASTRUCTURES

The PulseAudio daemon is deprecated

The PulseAudio daemon, and its packages **pulseaudio** and **alsa-plugins-pulseaudio**, have been deprecated and will be removed in a future major release.

Note that the PulseAudio client libraries and tools are not deprecated, this change only impacts the audio daemon that runs on the system.

You can use the PipeWire audio system as a replacement, which has also been the default audio daemon since RHEL 9.0. PipeWire also provides an implementation of the PulseAudio APIs.

Jira:RHELDOCS-19080[1]

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]

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.

Jira:RHELPLAN-157294[1]

9.17. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The mssql_accept_microsoft_odbc_driver_17_for_sql_server_eula variable has been deprecated

With a future major update of RHEL, the **mssql_accept_microsoft_odbc_driver_17_for_sql_server_eula** variable will no longer be supported in the **mssql** system role because the role can now install the **odbc** driver for **mssql_tools** version 17 and 18. Therefore, you must use the **mssql_accept_microsoft_odbc_driver_for_sql_server_eula** variable without the version number instead.

Important: If you use the deprecated variable with the version number **mssql_accept_microsoft_odbc_driver_17_for_sql_server_eula**, the role notifies you to use the new variable **mssql_accept_microsoft_odbc_driver_for_sql_server_eula**. However, the deprecated variable continues to work.

Jira:RHEL-69311

Deprecated variables in the podman RHEL system role: container_image_user and container_image_password

The **container_image_user** and **container_image_password** variables are deprecated. In a future major release of RHEL, these variables will be removed. You can use the **podman_registry_username** and **podman_registry_password** variables instead.

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

Jira:RHELDOCS-18803[1]

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.

Jira:RHELPLAN-95747[1]

9.18. VIRTUALIZATION

Specific IBM z16 CPU features have been deprecated.

With this update, the **te** and **cte** CPU features have been deprecated for IBM z16 KVM VMs. Note, however, that migrating a virtual machine with CPU model **host-model** from an IBM z16 host to an IBM z17 host does not require any adjustments to CPU feature settings.

Jira:RHEL-89415[1]

Live VM dumps have been deprecated

The **--live** option for the **virsh dump** command has become deprecated, and will be removed in a future release of RHEL. After the removal, if you attempt to create a virtual machine dump by using **virsh dump** with the **--live** option, the command will fail.

Jira:RHEL-57677

NIC device drivers related to iPXE are deprecated in RHEL 9

Internet Preboot eXecution Environment (iPXE) firmware provides a range of boot options over a network often used in environments, where machines need to boot remotely. Among others, it contains a large number of device drivers. The following have been marked as deprecated and will be removed in the RHEL 10 release:

- The complete **ipxe-roms** sub-RPM package
- Binary files containing device drivers from **ipxe-bootimgs-x86** sub-RPM package:
  - **/usr/share/ipxe/ipxe-i386.efi**
  - **/usr/share/ipxe/ipxe-x86_64.efi**
  - **/usr/share/ipxe/ipxe.dsk**
  - **/usr/share/ipxe/ipxe.iso**
  - **/usr/share/ipxe/ipxe.lkrn**
  - **/usr/share/ipxe/ipxe.usb**

Instead, iPXE now depends on the platform firmware to provide a NIC driver for the network boot. The **/usr/share/ipxe/ipxe-snponly-x86_64.efi** and **/usr/share/ipxe/undionly.kpxe** iPXE binary files are the part of the **ipxe-bootimgs** package and use the NIC driver provided by the platform firmware.

Jira:RHELDOCS-18531

libvirtd 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]

Using Windows Server 2012 or Windows 8 as a guest operating system is not supported

Because Microsoft ended support for the following versions of Windows, Red Hat also removed support for using these versions as a guest operating system in this update.

- Windows 8
- Windows 8.1

- Windows Server 2012
- Windows Server 2012 R2

Jira:RHEL-11810

Internal snapshots for VMs have been deprecated

Creating and reverting to a virtual machine (VM) snapshot has become deprecated for snapshots that use the *internal* snapshot mechanism, and will be removed in a future major release of RHEL. Instead, use snapshots with the *external* mechanism.

For more information, see Support limitations for virtual machine snapshots .

Jira:RHELDOCS-20135[1]

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]

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 SHA-2 algorithm, or later.

Jira:RHELPLAN-69533[1]

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 9.7.

Jira:RHELPLAN-81033[1]

qcow2-v2 image format is deprecated

With RHEL 9.7, 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 9.7 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.

Jira:RHELPLAN-75969[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'
```

Jira:RHELPLAN-114513[1]

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]

pmem device passthrough has become deprecated

With this update, the non-volatile memory library (**nvml**) packages have become deprecated, and will be removed in a future major version of RHEL. As a consequence, when the package removal occurs, it will no longer be possible to pass persistent memory (**pmem**) devices to the virtual machines (VMs). Note that emulated NVDIMM devices backed by volatile memory or files will still be available, but will not be possible to configure as persistent.

Jira:RHELDOCS-17989

Converting Xen virtual machines from RHEL 5 by using virt-v2v has been deprecated.

Using the **virt-v2v** tool to convert virtual machines from a RHEL 5 Xen host to KVM has become deprecated, and will be removed in a future major release of RHEL. For details, see the Red Hat Knowledge Base.

Jira:RHELDOCS-19193[1]

9.19. CONTAINERS

The rsyslog container image has been deprecated

The **rsyslog** container image has been deprecated and will be removed in a future major release.

Jira:RHELDOCS-19523[1]

The runc container runtime has been deprecated

The **runc** is deprecated and will be removed in RHEL 10.0. The default container runtime in RHEL 9 is crun. The crun is a fast and low-memory footprint OCI container runtime written in C. The crun binary is up to 50 times smaller and up to twice as fast as the runc binary. Using crun, you can also set a minimal number of processes when running your container. The crun runtime also supports OCI hooks.

Jira:RHEL-69742

The podman-tests package has been deprecated

The **podman-tests** package has been deprecated.

Jira:RHEL-67859

nodejs-18 and nodejs-18-minimal are deprecated

The **nodejs-18** and **nodejs-18-minimal** container images are now deprecated and will no longer receive feature updates. Use **nodejs-22** and **nodejs-22-minimal** instead.

Jira:RHELDOCS-20283[1]

The ruby-31 container image is deprecated

The **ruby-31** container image is deprecated and will no longer receive feature updates. Use the **ruby-33** container image instead.

Jira:RHELDOCS-20519[1]

php-81 container image is deprecated

The **php-81** container image is now deprecated and will no longer receive feature updates. Use **php-83** instead.

Jira:RHELDOCS-20718[1]

The Podman v5.0 deprecations

In RHEL 9.5, the following is deprecated in Podman v5.0:

- The system connections and farm information stored in the **containers.conf** file are now read-only. The system connections and farm information will now be stored in the **podman.connections.json** file, managed only by Podman. Podman continues to support the old configuration options such as **[engine.service_destinations]** and the **[farms]** section. You can still add connections or farms manually if needed; however, it is not possible to delete a connection from the **containers.conf** file with the **podman system connection rm** command.
- The **slirp4netns** network mode is deprecated and will be removed in a future major release of RHEL. The **pasta** network mode is the default network mode for rootless containers.
- The cgroups v1 for rootless containers is deprecated and will be removed in a future major release of RHEL.

Jira:RHELDOCS-19021[1]

The runc container runtime has been deprecated

The **runc** container runtime is deprecated and will be removed in a future major release of RHEL. The default container runtime is **crun**.

Jira:RHELDOCS-19012[1]

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.

Jira:RHELPLAN-117005[1]

rhel9/pause has been deprecated

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

Jira:RHELPLAN-127619[1]

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]

pasta as a network name has been deprecated

The support for **pasta** as a network name value is deprecated and will not be accepted in the next major release of Podman, version 5.0. You can use the **pasta** network name value to create a unique network mode within Podman by employing the **podman run --network** and **podman create --network** commands.

Jira:RHELDOCS-17038[1]

The BoltDB database backend has been deprecated

The BoltDB database backend is deprecated as of RHEL 9.4. In a future version of RHEL, the BoltDB database backend will be removed and will no longer be available to Podman. For Podman, use the SQLite database backend, which is now the default as of RHEL 9.4.

Jira:RHELDOCS-17495[1]

The CNI network stack has been deprecated

The Container Network Interface (CNI) network stack is deprecated and will be removed in a future release. Use the Netavark network stack instead. For more information, see Switching the network stack from CNI to Netavark.

Jira:RHELDOCS-17518[1]

The Podman v5.0 upcoming deprecations

The following will be deprecated in the upcoming Podman v5.0, which will be released in RHEL 9.5 and RHEL 10.0 Beta:

- The BoltDB database backend will be deprecated. The new SQLite database backend is available.
- The **containers.conf** file will be read-only. The system connections and farm information will be stored in the **podman.connections.json** file, managed only by Podman. Podman continues to support the old configuration options such as **[engine.service_destinations]** and the **[farms]** section. You can still add connections or farms manually if needed, however, it is not possible to delete a connection from the **containers.conf** file with the **podman system connection rm** command.

The following changes are planned for RHEL 10.0 Beta:

- The **pasta** network mode will be the default network mode for rootless containers. The **slirp4netns** network mode will be deprecated.
- The cgroupv1 will be deprecated.
- The CNI network stack will be deprecated.

Jira:RHELDOCS-17462[1]

The rhel9/openssl has been deprecated

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

Jira:RHELDOCS-18106[1]

9.20. DEPRECATED PACKAGES

This section lists packages that have been deprecated and will probably not be included in a future major release of Red Hat Enterprise Linux.

For changes to packages between RHEL 8 and RHEL 9, see Changes to packages in the *Considerations in adopting RHEL 9* document.

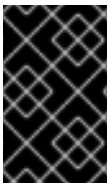

IMPORTANT

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:

- aacraid
- adwaita-gtk2-theme
- af_key
- anaconda-user-help

- aajohan-comfortaa-fonts
- adwaita-gtk2-theme
- adwaita-qt5
- anaconda-user-help
- ant-javamail
- apr-util-bdb
- aspnetcore-runtime-7.0
- aspnetcore-targeting-pack-6.0
- aspnetcore-targeting-pack-7.0
- atkmm
- atlas
- atlas-devel
- atlas-z14
- atlas-z15
- authselect-compat
- autoconf-latest
- autoconf271
- 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
- avahi-autoipd
- babl
- bacula-client
- bacula-common

- bacula-console
- bacula-director
- bacula-libs
- bacula-libs-sql
- bacula-logwatch
- bacula-storage
- bind9.18-libs
- bitmap-fangsongti-fonts
- bnx2
- bnx2fc
- bnx2i
- bogofilter
- Box2D
- brasero-nautilus
- cairomm
- cheese
- cheese-libs
- clucene-contribs-lib
- clucene-core
- clutter
- clutter-gst3
- clutter-gtk
- cnic
- cockpit-composer
- cogl
- compat-hesiod
- compat-locales-sap
- compat-locales-sap-common
- compat-openssl11

- compat-paratype-pt-sans-fonts-f33-f34
- compat-sap-c++-12
- compat-sap-c++-13
- containernetworking-plugins
- containers-common-extra
- culmus-aharoni-clm-fonts
- culmus-caladings-clm-fonts
- culmus-david-clm-fonts
- culmus-drugulin-clm-fonts
- culmus-ellinia-clm-fonts
- culmus-fonts-common
- culmus-frank-ruehl-clm-fonts
- culmus-hadasim-clm-fonts
- culmus-miriam-clm-fonts
- culmus-miriam-mono-clm-fonts
- culmus-nachlieli-clm-fonts
- culmus-simple-clm-fonts
- culmus-stamashkenaz-clm-fonts
- culmus-stamsefarad-clm-fonts
- culmus-yehuda-clm-fonts
- curl-minimal
- daxio
- dbus-glib
- dbus-glib-devel
- devhelp
- devhelp-libs
- dhcp-client
- dhcp-common
- dhcp-relay

- dhcp-server
- dotnet-apphost-pack-6.0
- dotnet-apphost-pack-7.0
- dotnet-hostfxr-6.0
- dotnet-hostfxr-7.0
- dotnet-runtime-6.0
- dotnet-runtime-7.0
- dotnet-sdk-6.0
- dotnet-sdk-7.0
- dotnet-targeting-pack-6.0
- dotnet-targeting-pack-7.0
- dotnet-templates-6.0
- dotnet-templates-7.0
- double-conversion
- efs-utils
- enchant
- enchant-devel
- eog
- evince
- evince-libs
- evince-nautilus
- evince-previewer
- evince-thumbnailer
- evolution
- evolution-bogofilter
- evolution-data-server-ui
- evolution-data-server-ui-devel
- evolution-devel
- evolution-ews

- evolution-ews-langpacks
- evolution-help
- evolution-langpacks
- evolution-mapi
- evolution-mapi-langpacks
- evolution-pst
- evolution-spamassassin
- festival
- festival-data
- festvox-slt-arctic-hts
- firefox
- firefox
- firefox-x11
- flite
- flite-devel
- fltk
- flute
- firewire-core
- fontawesome-fonts
- gc
- gcr-base
- gdisk
- 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
- ghc-srpm-macros
- ghostscript-x11
- git-p4
- gl-manpages
- glade
- glade-libs
- glibmm24
- gnome-backgrounds
- gnome-backgrounds-extras
- gnome-common
- gnome-logs
- gnome-photos
- gnome-photos-tests
- gnome-screenshot
- gnome-session-xsession
- gnome-shell-extension-panel-favorites

- gnome-shell-extension-updates-dialog
- gnome-terminal
- gnome-terminal-nautilus
- gnome-themes-extra
- gnome-tweaks
- gnome-video-effects
- google-noto-cjk-fonts-common
- google-noto-sans-cjk-ttc-fonts
- google-noto-sans-khmer-ui-fonts
- google-noto-sans-lao-ui-fonts
- google-noto-sans-thai-ui-fonts
- gspell
- gtksourceview4
- gtk2
- gtk2-devel
- gtk2-devel-docs
- gtk2-immodule-xim
- gtk2-immodules
- gtkmm30
- gtksourceview4
- gubbi-fonts
- gvfs-devel
- ha-openstack-support
- hexchat
- hesiod
- highcontrast-icon-theme
- http-parser
- ibus-gtk2
- initial-setup

- initial-setup-gui
- inkscape
- inkscape-docs
- inkscape-view
- iptables-devel
- iptables-libs
- iptables-nft
- iptables-nft-services
- iptables-utils
- iputils-ninfod
- ipxe-roms
- jakarta-activation2
- java-1.8.0-openjdk
- java-1.8.0-openjdk-demo
- java-1.8.0-openjdk-devel
- java-1.8.0-openjdk-headless
- java-1.8.0-openjdk-javadoc
- java-1.8.0-openjdk-javadoc-zip
- java-1.8.0-openjdk-src
- java-11-openjdk
- java-11-openjdk-demo
- java-11-openjdk-devel
- java-11-openjdk-headless
- java-11-openjdk-javadoc
- java-11-openjdk-javadoc-zip
- java-11-openjdk-jmods
- java-11-openjdk-src
- java-11-openjdk-static-libs
- java-17-openjdk

- java-17-openjdk-demo
- java-17-openjdk-devel
- java-17-openjdk-headless
- java-17-openjdk-javadoc
- java-17-openjdk-javadoc-zip
- java-17-openjdk-jmods
- java-17-openjdk-src
- java-17-openjdk-static-libs
- jboss-jaxrs-2.0-api
- jboss-logging
- jboss-logging-tools
- jdeparser
- jigawatts
- jigawatts-javadoc
- julietaula-montserrat-fonts
- kacst-art-fonts
- kacst-book-fonts
- kacst-decorative-fonts
- kacst-digital-fonts
- kacst-farsi-fonts
- kacst-fonts-common
- kacst-letter-fonts
- kacst-naskh-fonts
- kacst-office-fonts
- kacst-one-fonts
- kacst-pen-fonts
- kacst-poster-fonts
- kacst-qurn-fonts
- kacst-screen-fonts

- kacst-title-fonts
- kacst-titlel-fonts
- khmer-os-battambang-fonts
- khmer-os-bokor-fonts
- khmer-os-content-fonts
- khmer-os-fasthand-fonts
- khmer-os-freehand-fonts
- khmer-os-handwritten-fonts
- khmer-os-metal-chrieng-fonts
- khmer-os-muol-fonts
- khmer-os-muol-fonts-all
- khmer-os-muol-pali-fonts
- khmer-os-siemreap-fonts
- kmod-kvdo
- lasso
- libabw
- libadwaita-qt5
- libbase
- libblockdev-kbd
- libcanberra-gtk2
- libcdio-paranoia
- libcdio-paranoia-devel
- libcdr
- libcmis
- libdazzle
- libdb
- libdb-devel
- libdb-utils
- libdmx

- libepubgen
- libetonyek
- libexttextcat
- libfonts
- libformula
- libfreehand
- libgdata
- libgdata-devel
- libgnomekbd
- libiscsi
- libiscsi-utils
- liblangtag
- liblangtag-data
- liblayout
- libloader
- libmatchbox
- libmspub
- libmwaw
- libnsl2
- libnumbertext
- libodfgen
- liborcus
- libotr
- libpagemaker
- libpmem
- libpmem-debug
- libpmem-devel
- libpmem2
- libpmem2-debug

- libpmem2-devel
- libpmemblk
- libpmemblk-debug
- libpmemblk-devel
- libpmemlog
- libpmemlog-debug
- libpmemlog-devel
- libpmemobj
- libpmemobj++-devel
- libpmemobj++-doc
- libpmemobj-debug
- libpmemobj-devel
- libpmempool
- libpmempool-debug
- libpmempool-devel
- libpng15
- libpst-libs
- libqxp
- LibRaw
- 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-kn
- 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-voikko
- libreoffice-wiki-publisher
- libreoffice-writer
- libreoffice-x11
- libreoffice-xsltfilter
- libreofficekit
- libreport
- libreport-anaconda
- libreport-cli
- libreport-filesystem
- libreport-gtk
- libreport-plugin-bugzilla
- libreport-plugin-reportuploader
- libreport-rhel-anaconda-bugzilla
- libreport-web
- librepository
- librevenge
- librevenge-gdb

- libserializer
- libsigc++20
- libsigsegv
- libsmbios
- libsoup
- libsoup-devel
- libstaroffice
- libstemmer
- libstoragemgmt-smis-plugin
- libteam
- libuser
- libuser-devel
- libvisio
- libvisual
- libwpd
- libwpe
- libwpe-devel
- libwpg
- libwps
- libxcrypt-compat
- libxklavier
- libXp
- libXp-devel
- libXScrnSaver
- libXScrnSaver-devel
- libXxf86dga
- libXxf86dga-devel
- libzmf
- lklug-fonts

- lohit-gurmukhi-fonts
- lpsolve
- man-pages-overrides
- mcpp
- memkind
- mesa-libGLw
- mesa-libGLw-devel
- mlocate
- mod_auth_mellon
- mod_jk
- mod_security
- mod_security-mlogc
- mod_security_crs
- motif
- motif-devel
- mythes
- mythes-bg
- mythes-ca
- mythes-cs
- mythes-da
- mythes-de
- mythes-el
- mythes-en
- mythes-eo
- mythes-es
- mythes-fr
- mythes-ga
- mythes-hu
- mythes-it

- mythes-lv
- mythes-nb
- mythes-nl
- mythes-nn
- mythes-pl
- mythes-pt
- mythes-ro
- mythes-ru
- mythes-sk
- mythes-sl
- mythes-sv
- mythes-uk
- navilu-fonts
- nbdkit-gzip-filter
- neon
- NetworkManager-initscripts-updown
- nginx
- nginx-all-modules
- nginx-core
- nginx-filesystem
- nginx-mod-devel
- nginx-mod-http-image-filter
- nginx-mod-http-perl
- nginx-mod-http-xslt-filter
- nginx-mod-mail
- nginx-mod-stream
- nispor
- nscd
- nvme-stas

- opal-firmware
- opal-prd
- opal-prd
- opal-utils
- openal-soft
- openchange
- openscap-devel
- openscap-python3
- openslp-server
- overpass-fonts
- paktype-naqsh-fonts
- paktype-tehreer-fonts
- pam_ssh_agent_auth
- pangomm
- pentaho-libxml
- pentaho-reporting-flow-engine
- perl-AnyEvent
- perl-B-Hooks-EndOfScope
- perl-Class-Accessor
- perl-Class-Data-Inheritable
- perl-Class-Singleton
- perl-Class-Tiny
- perl-Crypt-OpenSSL-Bignum
- perl-Crypt-OpenSSL-Random
- perl-Crypt-OpenSSL-RSA
- perl-Date-ISO8601
- perl-DateTime
- perl-DateTime-Format-Builder
- perl-DateTime-Format-ISO8601

- perl-DateTime-Format-Strptime
- perl-DateTime-Locale
- perl-DateTime-TimeZone
- perl-DateTime-TimeZone-SystemV
- perl-DateTime-TimeZone-Tzfile
- perl-DB_File
- perl-Devel-CallChecker
- perl-Devel-Caller
- perl-Devel-LexAlias
- perl-Digest-SHA1
- perl-Dist-CheckConflicts
- perl-DynaLoader-Functions
- perl-Encode-Detect
- perl-Eval-Closure
- perl-Exception-Class
- perl-File-chdir
- perl-File-Copy-Recursive
- perl-File-Find-Object
- perl-File-Find-Rule
- perl-HTML-Tree
- perl-Importer
- perl-Mail-AuthenticationResults
- perl-Mail-DKIM
- perl-Mail-Sender
- perl-Mail-SPF
- perl-MIME-Types
- perl-Module-Implementation
- perl-Module-Pluggable
- perl-namespace-autoclean

- perl-namespace-clean
- perl-Net-CIDR-Lite
- perl-Net-DNS
- perl-NetAddr-IP
- perl-Number-Compare
- perl-Package-Stash
- perl-Package-Stash-XS
- perl-PadWalker
- perl-Params-Classify
- perl-Params-Validate
- perl-Params-ValidationCompiler
- perl-Perl-Destruct-Level
- perl-Ref-Util
- perl-Ref-Util-XS
- perl-Scope-Guard
- perl-Specio
- perl-Sub-Identify
- perl-Sub-Info
- perl-Sub-Name
- perl-Switch
- perl-Sys-CPU
- perl-Sys-MemInfo
- perl-Test-LongString
- perl-Test-Taint
- perl-Variable-Magic
- perl-XML-DOM
- perl-XML-RegExp
- perl-XML-Twig
- pinfo

- pki-jackson-annotations
- pki-jackson-core
- pki-jackson-databind
- pki-jackson-jaxrs-json-provider
- pki-jackson-jaxrs-providers
- pki-jackson-module-jaxb-annotations
- pki-resteasy-client
- pki-resteasy-core
- pki-resteasy-jackson2-provider
- pki-resteasy-servlet-initializer
- plymouth-theme-charge
- pmdk-convert
- pmempool
- podman-plugins
- poppler-qt5
- postgresql-test-rpm-macros
- power-profiles-daemon
- pulseaudio-module-x11
- python-botocore
- python-gflags
- python-netifaces
- python-pyroute2
- python-qt5-rpm-macros
- python3-bind
- python3-chardet
- python3-lasso
- python3-libproxy
- python3-libreport
- python3-netifaces

- python3-nispor
- python3-py
- python3-pycdlib
- python3-pycurl
- python3-pyqt5-sip
- python3-pyrsistent
- python3-pysocks
- python3-pytz
- python3-pywbem
- python3-qt5
- python3-qt5-base
- python3-requests+security
- python3-requests+socks
- python3-scour
- python3-toml
- python3-tomli
- python3-tracer
- python3-wx-siplib
- python3.11
- python3.11-cffi
- python3.11-charset-normalizer
- python3.11-cryptography
- python3.11-devel
- python3.11-idna
- python3.11-libs
- python3.11-lxml
- python3.11-mod_wsgi
- python3.11-numpy
- python3.11-numpy-f2py

- python3.11-pip
- python3.11-pip-wheel
- python3.11-ply
- python3.11-psycopg2
- python3.11-pycparser
- python3.11-PyMySQL
- python3.11-PyMySQL+rsa
- python3.11-pysocks
- python3.11-pyyaml
- python3.11-requests
- python3.11-requests+security
- python3.11-requests+socks
- python3.11-scipy
- python3.11-setuptools
- python3.11-setuptools-wheel
- python3.11-six
- python3.11-tkinter
- python3.11-urllib3
- python3.11-wheel
- python3.12-PyMySQL+rsa
- qgnomeplatform
- qla4xxx
- 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
- raptor2
- rasqal
- redis
- redis-devel
- redis-doc
- redland
- rpmlint
- runc
- saab-fonts
- sac
- satyr
- scap-workbench
- SDL2
- sendmail
- sendmail-cf
- sendmail-doc
- setxkbmap
- sgabios
- sgabios-bin
- sil-scheherazade-fonts
- spamassassin

- speech-tools-libs
- suitesparse
- sushi
- team
- teamd
- texlive-xdvi
- thai-scalable-fonts-common
- thai-scalable-garuda-fonts
- thai-scalable-kinnari-fonts
- thai-scalable-loma-fonts
- thai-scalable-norasi-fonts
- thai-scalable-purisa-fonts
- thai-scalable-sawasdee-fonts
- thai-scalable-tlwgmono-fonts
- thai-scalable-tlwgtypewriter-fonts
- thai-scalable-tlwgtypist-fonts
- thai-scalable-tlwgtypo-fonts
- thai-scalable-umpush-fonts
- thunderbird
- tigervnc
- tigervnc-icons
- tigervnc-license
- tigervnc-selinux
- tigervnc-server
- tigervnc-server-minimal
- tigervnc-server-module
- totem-pl-parser
- tracer-common
- ucs-miscfixed-fonts

- usb_modeswitch
- usb_modeswitch-data
- usbredir-server
- usermode-gtk
- webkit2gtk3
- webkit2gtk3-devel
- webkit2gtk3-jsc
- webkit2gtk3-jsc-devel
- wpebackend-fdo
- wpebackend-fdo-devel
- xmlrpc-c
- xmlsec1-gcrypt
- xmlsec1-gcrypt-devel
- xmlsec1-gnutls
- xmlsec1-gnutls-devel
- xorg-x11-drivers
- xorg-x11-drv-dummy
- xorg-x11-drv-evdev
- xorg-x11-drv-fbdev
- xorg-x11-drv-libinput
- xorg-x11-drv-v4l
- xorg-x11-drv-vmware
- xorg-x11-drv-wacom
- xorg-x11-drv-wacom-serial-support
- xorg-x11-server-common
- xorg-x11-server-utils
- xorg-x11-server-Xdmx
- xorg-x11-server-Xephyr
- xorg-x11-server-Xnest

- xorg-x11-server-Xorg
- xorg-x11-server-Xvfb
- xorg-x11-utils
- xorg-x11-xbitmaps
- xorg-x11-xinit
- xorg-x11-xinit-session
- xsane
- xsane-common
- xxhash
- xxhash-libs
- yajl
- yelp
- yelp-libs
- yp-tools
- ypbind
- ypserv
- zhongyi-song-fonts

CHAPTER 10. KNOWN ISSUES

This part describes known issues in Red Hat Enterprise Linux 9.7.

10.1. 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.

Workaround: Verify that the BaseOS and AppStream repositories are available to the installation program or use the **authselect** Kickstart command during installation.

Jira:RHELPLAN-10061[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.

Workaround: 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.

Jira:RHELPLAN-110940[1]

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 installation program 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.

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 .

Jira:RHELPLAN-53644[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.

Workaround: 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.

Workaround: 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.

Workaround: 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.

Jira:RHELPLAN-110191[1]

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

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

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

Workaround: You can use another filesystem for /**boot**, for example ext4.

Jira:RHELPLAN-94811[1]

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.
```

Workaround: Restart the installation process.

Jira:RHELPLAN-118420[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 installation program does not properly create some mount points for custom partitions. As a consequence, the installation stops with the following error:

```
The command 'mount --bind /mnt/sysimage/data /mnt/sysroot/data' exited with the code 32.
```

Workaround:

- 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.

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**.

Jira:RHELPLAN-130370[1]

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.

Workaround:

- Configure the network, for example using the **nmcli** tool, as a part of the **%pre** script.
- Use the installation program 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.

Jira:RHELPLAN-150080[1]

Images built with the stig profile remediation fail 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.

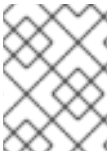

NOTE

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**.

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

Kickstart installation fails due to missing packages with systemd service files in %packages section

If the Kickstart file uses the **services --enabled=...** directive to enable **systemd** services and packages containing the specified service file are not included in the **%packages** section, the RHEL installation process fails with the following error:

```
Error enabling service <name_of_the_service>
```

Workaround: Include the respective package with the service file in Kickstart's **%packages** section. As a result, RHEL installation completes, enabling expected services during installation.

Jira:RHEL-9633[1]

Root filesystem are not expanded by default

When you use a base container image, that does not include **cloud-init** to create an AMI or QCOW2 container image by using **bootc-image-builder**, the root filesystem size is not expanded dynamically on boot to the full size of the provisioned virtual disk.

Workaround: Apply one of the following available options:

- Include **cloud-init** in the image.
- Include custom logic in the container image to expand the root filesystem, for example:

```
/usr/bin/growpart /dev/vda 4
unshare -m bin/sh -c 'mount -o remount,rw /sysroot && xfs_growfs /sysroot'
```

- Include a custom logic to use the additional space for secondary filesystems, for example, **/var/lib/containers**.

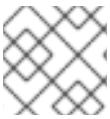

NOTE

By default, the physical root storage is mounted at the /**sysroot** partition.

Jira:RHEL-33208

Unable to build ISOs from a signed container

Trying to build an ISO disk image from a GPG or a simple signed container results in an error, similar to the following:

```
manifest - failed
Failed
Error: cannot run osbuild: running osbuild failed: exit status 1
```

```
2024/04/23 10:56:48 error: cannot run osbuild: running osbuild failed: exit status 1
```

This happens because the system fails to get the image source signatures.

Workaround: You can either remove the signature from the container image or build a derived container image. For example, to remove the signature, you can run the following command:

```
$ sudo skopeo copy --remove-signatures containers-storage:registry.redhat.io/rhel9/rhel-bootc:9.4
containers-storage:registry.redhat.io/rhel9/rhel-bootc:9.4
$ sudo podman run \
      --rm \
      -it \
      --privileged \
      --pull=newer \
      --security-opt label=type:unconfined_t \
      -v /var/lib/containers/storage:/var/lib/containers/storage \
      -v ~/images/iso:/output \
      quay.io/centos-bootc/bootc-image-builder \
      --type iso --local \
      registry.redhat.io/rhel9/rhel-bootc:9.4
```

To build a derived container image, and avoid adding a simple GPG signatures to it, see the Signing container images product documentation.

Jira:RHEL-34807

RHEL images on Azure marked as LVM require default layout resizing

When using **system-reinstall-bootc** or **bootc install** on Azure, RHEL images marked as LVM will require resizing the default layout.

Workaround: Use RHEL images labeled as RAW. This does not require resizing the default layout.

Jira:RHELDOCS-19945[1]

bootc-image-builder does not support building images from private registries

Currently, you cannot build base disk images which come from private registries by using **bootc-image-builder**.

Workaround: Copy the private registry into your localhost, then build the image with the following arguments:

- **--local**
- **localhost/<image name>:tag** as the image

For example, to build your image:

```
sudo podman run \
--rm \
-it \
--privileged \
--pull=newer \
--security-opt label=type:unconfined_t \
-v ./config.toml:/config.toml \
```

```
-v ./output:/output \
-v /var/lib/containers/storage:/var/lib/containers/storage \
registry.redhat.io/rhel9/bootc-image-builder:latest
--type qcow2 \
--local \
quay.io/<namespace>/<image>:<tag>
```

Jira:RHELDOCS-18720[1]

SELinux autorelabel in the Rescue Mode might cause reboot loop

Accessing a file system in the **rescue** mode triggers SELinux to autorelabel the file system on the next boot, which continues until SELinux runs in the **permissive** mode. Consequently, the system might go into an infinite loop of reboots after exiting the **rescue** mode as it cannot delete the **/.autorelabel** file.

Workaround: Switch to the **permissive** mode by adding **enforcing=0** to the kernel command line on the next boot. The system displays a warning message as a preventive measure that informs about the possibility of this issue when accessing the file system in the **rescue** mode.

Jira:RHEL-14005

Hostname resolution fails with encrypted DNS and custom CA in boot options

While using the **inst.repo=** or **inst.stage2=** boot options in the kernel command line along with a remote installation URL, an encrypted DNS, and a custom CA certificate in the kickstart file, the installation program attempts to download the **install.img** stage2 image before processing the kickstart file. Consequently, the hostname resolution fails, leading to display of some errors before successfully fetching the stage2 image.

Workaround: Define the installation source in the kickstart file instead of the kernel command line.

Jira:RHEL-80867

Bonding device with LACP takes longer to become operational, causing subscription failures

When configuring a bonding device with LACP by using both kernel command-line boot options and a Kickstart file, the connection is created during the **initramfs** stage but reactivated in Anaconda. As a consequence, it causes a temporary disruption that leads to system subscription failure via the **rhsm** Kickstart command.

Workaround: Add **--no-activate** to the Kickstart network configuration to keep the network operational. As a result, the system subscription completes successfully.

Jira:RHELDOCS-19852[1]

The services Kickstart command fails to disable the firewalld service

A bug in Anaconda prevents the **services --disabled=firewalld** command from disabling the **firewalld** service in Kickstart.

Workaround: Use the **firewall --disabled** command instead. As a result, the **firewalld** service is disabled properly.

Jira:RHEL-82566

Kickstart installation fails with an unknown disk error when **ignoredisk** command precedes **iscsi** command

Installing RHEL by using the kickstart method fails if the **ignoredisk** command is placed before the **iscsi** command. This issue occurs because the **iscsi** command attaches the specified iSCSI device during command parsing, while the **ignoredisk** command resolves device specifications simultaneously. If the **ignoredisk** command references an iSCSI device name before it is attached by the **iscsi** command, the installation fails with an "unknown disk" error.

Workaround: Ensure that the **iscsi** command is placed before the **ignoredisk** command in the Kickstart file to reference the iSCSI disk and enable successful installation.

Jira:RHEL-13837

Installation program fails if /boot partition is not created when using **ostreecontainer**

When using the **ostreecontainer** Kickstart command to install a bootable container, the installation fails if the **/boot** partition is not created. This issue occurs because the installation program requires a dedicated **/boot** partition to proceed with the container deployment.

Workaround: Ensure that a **/boot** partition is defined in the Kickstart file or manually created during the installation process.

Jira:RHEL-66155

Insufficient disk space can cause deployment failure

Deploying a bootc container image on a package mode system without enough free disk space can result in installation errors and prevent the system from booting. Ensure adequate disk space is available for the image to install and adjust the provision logical volume before deployment.

Jira:RHELDOCS-19948[1]

Anaconda may not work correctly on **s390x** and **ppc64le** architectures

Image mode for RHEL supports **pp64le** and **s390x** architectures besides the already supported **x86_64** and ARM architectures. However, Anaconda may not function correctly on s390x and ppc64le architectures.

Jira:RHELDOCS-19496[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.

Jira:RHELDOCS-20471[1]

bootc installation now checks for sufficient disk space

Previously, the bootc installation failed due to insufficient free space during installation, causing the system to fail to boot. With this update, bootc installation now checks for sufficient disk space before overwriting the host root filesystem. As a result, the failure to allocate space during bootc installation is

resolved, ensuring the system boots correctly.

Jira:RHEL-90381

10.2. 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.

Workaround: Configure servers and clients to use TLS version 1.2 as the highest TLS protocol version available.

Jira:RHELPLAN-50959[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.

Workaround: 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.

Jira:RHELPLAN-48241[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

Workaround: 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**

Jira:RHELPLAN-113842[1]

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.
```

Workaround: 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
   ```

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

   ```
   # ANSIBLE_COLLECTIONS_PATH=/usr/share/rhc-worker-playbook/ansible/collections/ansible_collections/ ansible-playbook -c local -i localhost, rhel9-playbook-cis_server_l1.yml
   ```

   Replace ***cis_server_l1*** with the ID of the profile against which you want to remediate the system.

As a result, the Ansible content is processed correctly.

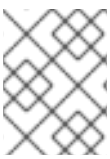

NOTE

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.

Workaround: 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 malformatted**.

Workaround: 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.

Workaround: 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]

Missing files in trustdb cause denials for fapolicyd

When **fapolicyd** is installed with the Ansible DISA STIG profile, a race condition causes the **trustdb** database to be out of sync with the **rpmdb** database. As a consequence, missing files in **trustdb** cause denials on the system.

Workaround: Restart **fapolicyd** or run the Ansible DISA STIG profile again.

Jira:RHEL-24345[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]

RPM packages signed with MLDSA-87 fail to install in FIPS mode

The post-quantum cryptography (PQC) algorithms are not FIPS-validated and are not available in the FIPS provider. This causes the import of MLDSA-87 PQC keys into the RPM database and PQC signature verification to fail in FIPS mode.

Workaround: Do not enable the DNF plugin to support PQC signatures in FIPS mode. As a result, the system verifies packages in FIPS mode through classical signatures.

Jira:RHEL-111478[1]

OpenSSL no longer creates X.509 v1 certificates

With the OpenSSL TLS toolkit 3.2.1 introduced in RHEL 9.5, you can no longer create certificates in the X.509 version 1 format using the **openssl** CA tool. The X.509 v1 format does not meet current web requirements.

Jira:RHEL-40605

OpenSSH no longer logs timeout before authentication

OpenSSH does not record a timeout before authentication for **$IP port $PORT** to the log. This might be important because the Fail2Ban intrusion prevention daemon and similar systems use these log records in its **mdre-ddos** regular expression and no longer ban the IPs of clients that attempt this type of attack. There is currently no known workaround for this problem.

Jira:RHEL-45727

Updating the NSS database password corrupts the ML-DSA seed

Generating an ML-DSA key begins with a seed, which is sufficient to derive the key. However, the keys can also be expanded to accelerate subsequent operations. If you have ML-DSA keys in an NSS database, either generated or imported, both the expanded format and the seed are likely stored. Due to a bug in how NSS handles database re-encryption, if you change the password of the database, the seed attribute is not updated to accommodate the new password, and its value is permanently lost, even with the knowledge of the previous password.

To work around this problem, export the key before updating the password and re-import it after the update.

Jira:RHEL-127671[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**.

Jira:RHELPLAN-115609[1]

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.

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.

Jira:RHELPLAN-107318[1]

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.

Workaround: 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.

Jira:RHELPLAN-117566[1]

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**

Workaround: See the related Knowledgebase article.

Jira:RHELPLAN-145263[1]

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.

Workaround: You can scan and remediate the system after the kickstart installation. This will fix the service-related issues.

Jira:RHELPLAN-44202[1]

Interoperability of FIPS:OSPP hosts impacted due to CNSA 1.0

The **OSPP** subpolicy has been aligned with Commercial National Security Algorithm (CNSA) 1.0. This affects the interoperability of hosts that use the **FIPS:OSPP** policy-subpolicy combination, with the following major aspects:

- Minimum RSA key size is mandated at 3072 bits.
- Algorithm negotiations no longer support AES-128 ciphers, the secp256r1 elliptic curve, and the FFDHE-2048 group.

Jira:RHEL-2735[1]

Missing rules in the SELinux policy block permissions to SQL databases

Missing permission rules from the SELinux policy block connections to SQL databases. Consequently, the FIDO Device Onboard (FDO) services **fdo-manufacturing-server.service**, **fdo-owner-onboarding-server.service**, and **fdo-rendezvous-server.service** cannot connect to FDO databases, such as PostgreSQL and SQLite. Therefore, the system cannot start the FDO by using the supported databases for credentials and other parameters, such as storing ownership vouchers.

Workaround: Perform the following steps:

1. Create a new file named **local_fdo_update.cil** and enter the missing SELinux policy rules:

```
(allow fdo_t etc_t (file (write)))
(allow fdo_t fdo_conf_t (file (append create rename setattr unlink write )))
(allow fdo_t fdo_var_lib_t (dir (add_name remove_name write )))
(allow fdo_t fdo_var_lib_t (file (create setattr unlink write )))
(allow fdo_t krb5_keytab_t (dir (search)))
(allow fdo_t postgresql_port_t (tcp_socket (name_connect)))
(allow fdo_t sssd_t (unix_stream_socket (connectto)))
(allow fdo_t sssd_var_run_t (sock_file (write)))
```

2. Install the policy module package:

```
# semodule -i local_fdo_update.cil
```

As a consequence, FDO can connect to the PostgreSQL database and also fix problems related to SQLite permissions over **/var/lib/fdo/**, where the SQLite database files are expected to be located.

Jira:RHEL-28814[1]

PQC for rpm-sequoia is always enabled in crypto-policies

In RHEL 10.1, the **rpm-sequoia** fails to verify dual-signed RPM packages if one of the algorithms used for signing is disabled in system-wide cryptographic policies. This problem is common on systems that have post-quantum (PQ) algorithms disabled and cannot install packages signed with both classic and PQ cryptography.

To prevent breaking the system, the enablement of PQ algorithms for **rpm-sequoia** is hardcoded on the **crypto-policies** level. As a result, PQ algorithms for **rpm-sequoia** are enabled regardless of any settings in **crypto-policies**.

Jira:RHEL-112697

10.3. SOFTWARE MANAGEMENT

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.

Workaround: 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**.

Jira:RHELPLAN-112860[1]

A security DNF upgrade fails for packages that change their architecture through the upgrade

The patch for BZ#2108969, released with the RHBA-2022:8295 advisory, introduced the following regression: The DNF upgrade using security filters fails for packages that change their architecture from or to **noarch** through the upgrade. Consequently, it can leave the system in a vulnerable state.

To work around this problem, perform the regular upgrade without security filters.

Jira:RHELPLAN-128381[1]

10.4. SHELLS AND COMMAND-LINE TOOLS

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.

Workaround: 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.

Workaround: 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.

Workaround: Configure firewalls to block or filter traffic on UDP and TCP port 427.

Jira:RHEL-6995[1]

The ReaR rescue image on UEFI systems with Secure Boot enabled fails to boot with the default settings

ReaR image creation by using the **rear mkrescue** or **rear mkbackup** command fails with the following message:

```
grub2-mkstandalone may fail to make a bootable EFI image of GRUB2 (no /usr/*/grub*/x86_64-efi/moddep.lst file)
(...)
grub2-mkstandalone: error: /usr/lib/grub/x86_64-efi/modinfo.sh doesn't exist. Please specify --target or --directory.
```

The missing files are part of the **grub2-efi-x64-modules** package. If you install this package, the rescue image is created successfully without any errors. When the **UEFI** Secure Boot is enabled, the rescue image is not bootable because it uses a boot loader that is not signed.

Workaround: Add the following variables to the **/etc/rear/local.conf** or **/etc/rear/site.conf** ReaR configuration file):

```
UEFI_BOOTLOADER=/boot/efi/EFI/redhat/grubx64.efi
```

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

With the suggested workaround, the image can be produced successfully even on systems without the **grub2-efi-x64-modules** package, and it is bootable on systems with Secure Boot enabled. In addition, during the system recovery, the bootloader of the recovered system is set to the **EFI** shim bootloader.

For more information about **UEFI**, **Secure Boot**, and **shim bootloader**, see the UEFI: what happens when booting the system Knowledge Base article.

Jira:RHELDOCS-18064[1]

The %util column produced by sar and iostat utilities is invalid

When you collect system usage statistics by using the **sar** or **iostat** utilities, the **%util** column produced by **sar** or **iostat** might contain invalid data.

Jira:RHEL-26275[1]

The lsb-release binary is not available in RHEL 9

The information in **/etc/os-release** was previously available by calling the **lsb-release** binary. This binary was included in the **redhat-lsb package**, which was removed in RHEL 9. Now, you can display information about the operating system, such as the distribution, version, code name, and associated metadata, by reading the **/etc/os-release** file. This file is provided by Red Hat and any changes to it will be overwritten with each update of the **redhat-release** package. The format of the file is **KEY=VALUE**, and you can safely source the data for a shell script.

Jira:RHELDOCS-16427[1]

10.5. INFRASTRUCTURE SERVICES

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.

Jira:RHELPLAN-122108[1]

Using the incorrect Perl database driver for MariaDB and MySQL can lead to unexpected results

The MariaDB database is a fork of MySQL. Over time, these services developed independently and are no longer fully compatible. These differences also affect the Perl database drivers. Consequently, if you use the **DBD::mysql** driver in a Perl application to connect to a MariaDB database, or the **DBD::MariaDB** driver to connect to a MySQL database, operations can lead to unexpected results. For example, the driver can return incorrect data from read operations. To avoid such problems, use the Perl driver in your application that matches the database service.

Red Hat only supports the following scenarios:

- The Perl **DBD::MariaDB** driver with a MariaDB database
- The Perl **DBD::mysql** driver with a MySQL database

Note that RHEL 8 contained only the **DBD::mysql** driver. If you plan to upgrade to RHEL 9 and then to RHEL 10 and your application uses a MariaDB database, install the **perl-DBD-MariaDB** package after the upgrade and modify your application to use the **DBD::MariaDB** driver.

For further details, see the Red Hat Knowledgebase solution Support of MariaDB/MySQL cross-database connection from Perl db drivers.

Jira:RHELDOCS-19728[1]

Hot-plugged memory is not available to VMs running on IBM Z by default

RHEL provides default udev rules that automatically configure memory onlining when you hot plug memory to virtual machines (VMs) with **virtio-mem**. However, current udev rules do not include VMs running on IBM Z. As a consequence, after hot-plugging memory to VMs running on IBM Z with **virtio-mem**, the memory is not immediately available in the VM.

To work around this problem, set the **memhp_default_state=online** kernel parameter in the VM and reboot it. For example:

```
# grubby --update-kernel=ALL --args=memhp_default_state=online
```

As a result, the hot-plugged memory is available in the VM.

Jira:RHEL-92781

10.6. NETWORKING

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.

Workaround: 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.

Jira:RHELPLAN-96004[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.

Workaround: Disable kTLS. As a result, with the workaround, it is possible to successfully update the session key.

Jira:RHELPLAN-99859[1]

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.

Workaround: 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.

Jira:RHELPLAN-100926[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.

Workaround: 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.

Jira:RHELPLAN-121205[1]

The iwl7260-firmware causes Wi-Fi issues on Intel Wi-Fi 6 AX200, AX210, and Lenovo ThinkPad P1 Gen 4

If you update the **iwl7260-firmware** or **iwl7260-wifi** driver to the version provided with RHEL 9.1 or later, the hardware might enter in an incorrect state and report its status incorrectly. Consequently, Intel Wi-Fi 6 cards might fail to function properly and display the following 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
```

Workaround: An unconfirmed workaround is to power off the system completely and then power it back on. Do not perform a reboot.

Jira:RHELPLAN-134771[1]

Issues in DPLL stability during PF resets

The Digital Phase-Locked Loop (DPLL) system experienced several issues, including uninitialized mutex usage and incorrect handling of pin phase adjustments, particularly during Physical Function (PF) resets. These issues led to unstable management of DPLL and pin configurations, causing inconsistent data states and connection mismanagement.

Workaround: To resolve this, mutexes were properly initialized, and mechanisms for updating pin phase adjustments, DPLL data, and connection states during PF resets were corrected. As a result, the DPLL system now performs reliably during resets, with accurate phase adjustments and consistent connection states, improving the overall stability of clock synchronization.

Jira:RHEL-36283[1]

10.7. KERNEL

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
```

Workaround: 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
```

Jira:RHELPLAN-147783[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.

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.

Jira:RHELPLAN-153123[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 **iotop** application does not show the **SWAPIN** and **IO%** statistics columns and displays the following warning:

```
CONFIG_TASK_DELAY_ACCT not enabled in kernel, cannot determine SWAPIN and IO%
```

The **Delay Accounting** functionality, using the **taskstats** interface, provides the delay statistics for all

tasks or threads that belong to a thread group. Delays in task execution occur when they wait for a kernel resource to become available, for example, a task waiting for a free CPU to run on. The statistics help in setting a task's CPU priority, I/O priority, and **rss** limit values appropriately.

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.

Jira:RHELPLAN-135779[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.

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.

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
```

Jira:RHELPLAN-115732[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.

Workaround: Build or put the extra modules against the latest stock kernel before you install the new kernel.

Jira:RHELPLAN-126922[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]

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.

Workaround: 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]

Kernel panic is encountered on IBM Power systems (ppc64le) when io_uring is enabled

In some cases,**ppc64le** systems encounter a kernel panic when using the **io_uring** kernel parameter due to intensive input-output operations. As a consequence, **ppc64le** stops working and requires a system restart. The data might get lost during the crash.

Workaround: Disable the **io_uring** feature by adding the following kernel parameter at boot time:

module.builtin=io_uring=0

Jira:RHEL-28702[1]

10.8. FILE SYSTEMS AND STORAGE

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.

Jira:RHELPLAN-105944[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.

Workaround: 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.

Jira:RHELPLAN-99108[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.

Workaround: Disable quota accounting by remounting the filesystem, with the quota option removed.

Jira:RHELPLAN-145001[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.

Jira:RHELPLAN-154195[1]

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.

Workaround: 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]

ARM-based systems fail to update with a 64k page size kernel when vdo is installed

While installing the **vdo** package, RHEL installs the **kmod-kvdo** package and a kernel with **4k** page size as dependencies. As a consequence, updates from RHEL 9.3 to 9.x fail because **kmod-kvdo** conflicts with the 64k kernel.

Workaround: Remove the **vdo** package and its dependencies before attempting to update.

Jira:RHEL-8354

lldpad is auto enabled even for qedf adapters

When using a QLogic Corp. FastLinQ QL45000 Series 10/25/40/50GbE, FCOE Controller automatically enables the **lldpad** daemon on systems running RHV. As a consequence, I/O operations are stopped with an error, for example, **[qedf_eh_abort:xxxx]:1: Aborting io_req=ff5d85a9dcf3xxxx**.

Workaround: DisableLink Layer Discovery Protocol (LLDP) and then enable it for interfaces that can be set on the **vdsm** configuration level. For more information, https://access.redhat.com/solutions/6963195.

Jira:RHEL-8104[1]

System fails to boot when iommu is enabled

By enabling the Input-Output Memory Management Unit (IOMMU) on AMD platforms when the BNX2I adapter is in use, a system fails to boot with the Direct Memory Access Remapping (DMAR) timeout errors.

Workaround: Disable the IOMMU before booting by using the kernel command-line option, **iommu=off**. As a result, the system boots without any errors.

Jira:RHEL-25730[1]

10.9. HIGH AVAILABILITY AND CLUSTERS

Removing duplicate route entries for IPv6 addresses in an IPsrcaddr resource

In Red Hat Enterprise Linux 9.4 and earlier, when you specified an IPv6 address for an **IPsrcaddr** resource, the **IPsrcaddr** resource agent created a duplicate route with a different metric when the metric was used for the subnet. For example, this happened when NetworkManager created another IP address on the IPv6 subnet. In this situation, the **IPsrcaddr** resource failed to start because there was more than one match for the IP address. As of Red Hat Enterprise Linux 9.5, the **IPsrcaddr** resource agent specifies the metric of an existing route when it is available and a second route is not created. If, however, you created an **IPaddr2** IPv6 resource that uses an IPv6 address before this upgrade, you must reboot your system to remove the duplicate route entry.

Jira:RHEL-32265[1]

10.10. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

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.

Jira:RHELPLAN-112043[1]

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.

Jira:RHELPLAN-143480[1]

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 **mariadbd** 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 **mariadbd** server daemon does not start.

- If you set **--ssl-fips-mode** to **OFF** on a FIPS-enabled system, the **mysqld** or **mariadbd** 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 .

Jira:RHELPLAN-92864[1]

Git fails to clone or fetch from repositories with potentially unsafe ownership

To prevent remote code execution and mitigate CVE-2024-32004, stricter ownership checks have been introduced in **Git** for cloning local repositories. With this update, **Git** treats local repositories with potentially unsafe ownership as dubious.

As a consequence, if you attempt to clone from a repository locally hosted through **git-daemon** and you are not the owner of the repository, **Git** returns a security alert about dubious ownership and fails to clone or fetch from the repository.

Workaround: Explicitly mark the repository as safe by executing the following command:

```
git config --global --add safe.directory /path/to/repository
```

Jira:RHELDOCS-18435[1]

10.11. COMPILERS AND DEVELOPMENT TOOLS

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.

Jira:RHELPLAN-117492[1]

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.

Workaround: 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.

Jira:RHELPLAN-90604[1]

10.12. IDENTITY MANAGEMENT

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.

Workaround: Enable support for the SHA-1 algorithm on your RHEL 9 systems with the following command:

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

Jira:RHELPLAN-114497[1]

The PKINIT authentication of a user fails if a RHEL 9 Kerberos agent communicates with a non-RHEL-9, 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.

Workaround: 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
```

Jira:RHELPLAN-113281[1]

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]

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:'
```

For more information, see the AD Domain Users unable to login in to the FIPS-compliant environment KCS solution.

Jira:RHEL-4888

The online backup and the online automembership rebuild tasks can acquire two locks resulting in a deadlock

If the online backup and the online automembership rebuild tasks attempt to acquire the same two locks in the opposite order, it can lead to an unrecoverable deadlock that requires you to stop and restart the server. To work around this problem, do not launch the online backup and the online automembership rebuild tasks in parallel.

Jira:RHELDOCS-18065[1]

Installing a RHEL 7 IdM client with a RHEL 9.2 and later 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.

Workaround: If upgrading the host to RHEL 8 before installing an IdM client on it is not an option, remove 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

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-autentication 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.

Workaround: 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:RHELDOCS-19846[1]

10.13. SSSD

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/sssd/sssd.conf** file. The default behavior is planned to be changed in a future release of RHEL.

Jira:RHELPLAN-155168[1]

SSSD retrieves incomplete list of members if the group size exceeds 1500 members

During the integration of SSSD with Active Directory, SSSD retrieves incomplete group member lists when the group size exceeds 1500 members. This issue occurs because Active Directory's MaxValRange policy, which restricts the number of members retrievable in a single query, is set to 1500 by default.

Workaround: Change the MaxValRange setting in Active Directory to accommodate larger group sizes.

Jira:RHELDOCS-19603[1]

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.

Workaround: 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.

Jira:RHELPLAN-44204[1]

10.14. 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.

Workaround: 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.

Jira:RHELPLAN-114314[1]

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.

Workaround: 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]

xorg -configure fails to create an Xorg configuration file on a virtual machine

Running **xorg -configure** to create the Xorg configuration file on virtual machines fails due to the lack of devices to configure. This issue leads to a configuration failure. To work around this issue, construct the **xorg.conf** file manually according to the guidelines stated in Xorg documentation, or use alternative mechanisms such as an Extended Display Identification Data (EDID) override to tweak display resolutions. With this workaround, the Xorg server functions with the correct configuration.

Jira:RHELDOCS-20196[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

10.15. 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]

10.16. THE WEB CONSOLE

VNC console in the RHEL web console does not work correctly on ARM64

Currently, when you import a virtual machine (VM) in the RHEL web console on ARM64 architecture and then you try to interact with it in the VNC console, the console does not react to your input.

Additionally, when you create a VM in the web console on ARM64 architecture, the VNC console does not display the last lines of your input.

Jira:RHEL-31993[1]

10.17. 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.

Workaround: Unmask the **firewalld.service**:

```
systemctl unmask firewalld.service
```

Jira:RHELPLAN-133165[1]

PostgreSQL, MariaDB, and MySQL do not work with RHEL in image mode

The PostgreSQL, MariaDB, and MySQL database management systems do not use the **sysusers.d** directories to populate users and working directories. MariaDB and MySQL also do not use the **tmpfiles.d** directory. As a consequence, the database user can be missing and the database systems are not able to initialize because their working directory is missing. There is currently no workaround for this issue.

Jira:RHELDOCS-21366[1]

Unable to register systems with environment names

The **rhc** system role fails to register the system when specifying environment names in **rhc_environment**.

Workaround: Use environment IDs instead of environment names while registering.

Jira:RHEL-1172

Running Microsoft SQL Server 2022 in high-availability mode as an SELinux-confined application does not work

Microsoft SQL Server 2022 on RHEL 9.4 and later supports running as an SELinux-confined application. However, due to a limitation in Microsoft SQL Server, running the service as an SELinux-confined application does not work in high-availability mode.

Workaround: You can run Microsoft SQL Server as an unconfined application if you require the service to be high available.

Note that this limitation also impacts installing Microsoft SQL Server when you use the **mssql** RHEL System Role to install this service.

Jira:RHELDOCS-17719[1]

10.18. 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.

Workaround: 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.

Jira:RHELPLAN-99854[1]

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

There is currently no workaround for this issue.

Jira:RHELPLAN-117234[1]

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. Regenerate initramfs. To do so, use the following steps in the VM and replace ***<kernelVersion>*** with the full version of the kernel that you want to rebuild.
   a. Back up the current **initramfs** configuration:

```
# cp /boot/initramfs-<kernelVersion>.img /boot/initramfs-<kernelVersion>.img.bak
```

   b. Build **initramfs**:

```
# dracut -f /boot/initramfs-<kernelVersion>.img  <kernelVersion>
```

3. Reboot the VM to verify successful boot.

Jira:RHELPLAN-114103[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.

Workaround: Manually turn on **erms** and **fsrm** in the BIOS of your host.

Jira:RHELPLAN-119655[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. There is currently no workaround for this issue.

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.

Workaround: 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 a VMs that uses virtual SR-IOV software, the network of the VM is not visible to other devices and the VM cannot be reached by commands such as **ping**. After the migration is finished, however, the problem no longer occurs.

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. There is currently no workaround for this issue.

Jira:RHELPLAN-97394[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.

Workaround: See the following Knowledgebase article: Migrated IdM users unable to log in due to mismatching domain SIDs.

Jira:RHELPLAN-109613[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.

Workaround: 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. There is currently no workaround for this issue.

Jira:RHELPLAN-63771[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 NVDIA 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 KnoweldgeBase .

Jira:RHELPLAN-141042[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.

Workaround: 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**

Workaround: 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
```

Workaround: Do not use AMD EPYC CPUs for NUMA node configurations.

Jira:RHELPLAN-150884[1]

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 .

Jira:RHELPLAN-118495[1]

virsh blkiotune --weight command fails to set the correct cgroup I/O controller value

Currently, using the **virsh blkiotune --weight** command to set the VM weight does not work as expected. The command fails to set the correct **io.bfq.weight** value in the cgroup I/O controller interface file. There is no workaround at this time.

Jira:RHELPLAN-83423[1]

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]

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. There is currently no workaround for this issue.

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.

Workaround: Boot the VM with the **multi_queue** option enabled.

Jira:RHEL-1084[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]

The virtio balloon driver sometimes does not work on Windows 10 and Windows 11 VMs

Under certain circumstances, the **virtio-balloon** driver does not work correctly on virtual machines (VMs) that use a Windows 10 or Windows 11 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. There is currently no workaround for this issue.

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. There is currently no workaround for this issue.

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
```

There is currently no workaround for this issue.

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.

Workaround: Update the **virtio** drivers by using Windows Device Manager.

Jira:RHEL-574[1]

TX queue size cannot be changed in VMs that use vhost-kernel

Currently, you cannot set up TX queue size on KVM virtual machines (VMs) that use **vhost-kernel** as a back end for the **virtio** network driver. As a consequence, you can use only the default value of 256 for the TX queue, which might prevent you from optimizing your VM network throughput. There is currently no workaround for this issue.

Jira:RHEL-1138[1]

VMs incorrectly report the vulnerable status for spec_rstack_overflow parameter on the AMD EPYC model

When you boot a host, it does not detect any vulnerabilities in the **spec_rstack_overflow** parameter. After querying the parameter for logs, it displays the message:

```
# cat /sys/devices/system/cpu/vulnerabilities/spec_rstack_overflow
Mitigation: Safe RET
```

After booting a VM on the same host, the VM detects a vulnerability in the **spec_rstack_overflow** parameter. And when you query the parameter for logs, it displays the message:

```
# cat /sys/devices/system/cpu/vulnerabilities/spec_rstack_overflow
Vulnerable: Safe RET, no microcode
```

However, this is a false warning message, and you can ignore the status of the **/sys/devices/system/cpu/vulnerabilities/spec_rstack_overflow** file inside the VM.

Jira:RHEL-17614[1]

Link status shows up on VM, even when status is down of e1000e or igb model interface

Before booting the VM, set the status of Ethernet link **down** for the **e1000** or **igb** model network interface. Despite this, after the VM boots, the network interface keeps the **up** status, because when you set the status of Ethernet link **down** and then stop and re-start the VM, it is automatically set back to **up**. Consequently, the correct state of network interface is not maintained.

Workaround: Set the network interface status to **down** inside the VM by using command:

```
# ip link set dev eth0 down
```

Alternatively, you can try to remove and add this network interface again while the VM is running.

Jira:RHEL-21867

SeaBIOS cannot boot from a disk with 4096 bytes sector size

When using SeaBIOS to boot a virtual machine (VM) from a disk that uses logical or physical sector size of 4096 bytes, the boot disk is not displayed as available, and booting the VM fails. To boot a VM from such a disk, use UEFI instead of SeaBIOS.

Jira:RHEL-7110

Windows Server 2019 virtual machines crash on boot if using more than 128 cores per CPU

Virtual machines (VMs) that use a Windows Server 2019 guest operating system currently fail to boot when they are configured to use more than 128 cores for a single virtual CPU (vCPU). Instead of booting, the VM displays a stop error on a blue screen.

Workaround: Use fewer than 128 core per vCPU.

Jira:RHELDOCS-18863[1]

Windows VM with VBS and IOMMU device fails to boot

When you boot a Windows VM with Virtualization Based Security (VBS) enabled and an Input-Output Memory Management Unit (IOMMU) device by using the **qemu-kvm** utility, the booting sequence only shows the boot screen, resulting in an incomplete booting process.

Workaround: Ensure the VM domain XML is configured as below:

```
<features>
  <ioapic driver='qemu'/>
</features>
<devices>
```

```
<iommu model='intel'>
  <driver intremap='on' eim='off' aw_bits='48'/>
  <alias name='iommu0'/>
</iommu>
<memballoon model='virtio'>
  <alias name='balloon0'/>
  <address type='pci' domain='0x0000' bus='0x03' slot='0x00' function='0x0'/>
  <driver iommu='on' ats='on'/>
</memballoon>
</devices>
```

Otherwise, the Windows VM cannot boot.

Jira:RHEL-45585[1]

VMs with 5-level page merging and a lot of memory sometimes fail to start

VMs with the following configuration fail to boot if you set the **host-phys-bits-limit** parameter to **49** or more:

- The VM has more than 1 TB of assigned memory
- The VM uses the 5-level page merging feature
- The host uses System Management Mode (SMM) in its firmware

Instead, attempting to boot the VM fails with **ERROR: Out of aligned pages**.

Workaround: Set the **host-phys-bits-limit** parameter to 48 or less.

Jira:RHEL-82759

A virtual machine with a large amount of bootable data disks might fail to start

If you attempt to start a virtual machine (VM) with a large amount of bootable data disks, the VM might fail to boot with this error: **Something has gone seriously wrong: import_mok_state() failed: Volume Full**

Workaround: Decrease the number of bootable data disks and use one system disk. To ensure the system disk is first in the boot order, add **boot order=1** to the device definition of the system disk in the XML configuration. For example:

```
<disk type='file' device='disk'>
  <driver name='qemu' type='qcow2'/>
  <source file='/path/to/disk.qcow2'/>
  <target dev='vda' bus='virtio'/>
  <boot order='1'/>
</disk>
```

Set boot order only for the system disk.

Jira:RHEL-68418

Windows 2025 VM slows down if assigned with a large number of vCPU

When assigned with 32 or more vCPUs, Windows Server 2025 virtual machines (VMs) slow down on a Red Hat Enterprise Linux host. Consequently, a Windows VM may boot slowly or be stuck during boot when the VM is configured with a large number of vCPUs.

Workaround: You can use the workaround at your own risk. Boot VM with small number of vCPUs to disable plaformclock on Windows Server. In command prompt with administrator privileges, run:

```
bcdedit /set useplatformclock no
```

Then, shut down the VM and reconfigure it with the desired large number of vCPUs. Also make sure that the **hv-time** option is enabled before starting the large VM again.

Jira:RHEL-62742[1]

VMs with large memory cannot boot on SEV-SNP host with AMD Genoa CPUs

Currently, virtual machines (VMs) cannot boot on hosts that use a 4th Generation AMD EPYC processor (also known as Genoa) and have the AMD Secure Encrypted Virtualization with Secure Nested Paging (SEV-SNP) feature enabled. Instead of booting, a kernel panic occurs in the VM.

Jira:RHEL-32892[1]

Installing the VirtIO-Win bundle cannot be canceled

Currently, if you start the installation of **virtio-win** drivers from the VirtIO-Win installer bundle in a Windows guest operating system, clicking the **Cancel** button during the installation does not correctly stop it. The installer wizard interface displays a "Setup Failed" screen, but the drivers are installed and the IP address of the guest is reset.

Jira:RHEL-53962, Jira:RHEL-53965

Windows VM running on Sapphire Rapids CPU with hypervisor launch type set to auto might fail to boot when restarted

If you set the hypervisor launch type to **auto** in a Windows virtual machine (VM) running on a Sapphire Rapids CPU, the VM might fail to boot when it is restarted. For example, you can set the hypervisor launch type to **auto** by using the **bcdedit /set hypervisorlaunchtype Auto** command.

Workaround: Do not set the hypervisor launch type to **auto** in the Windows VM.

Jira:RHEL-67699[1]

Hot-plugging vCPUs and memory to Windows guests with VBS does not work

Currently, Windows Virtualization-based Security (VBS) is not compatible with hot-plugging CPU and memory resources. As a consequence, attempting to attach memory or vCPUs to a running Windows virtual machine (VM) with VBS enabled only adds the resources to the VM after the guest system is restarted.

Jira:RHEL-66229, Jira:RHELDOCS-19066

NetworkManager-wait-online.service fails to start on Azure VMs with Accelerated Networking

When you launch a Red Hat Enterprise Linux VM of Azure platform with the Accelerated Networking feature, also known as Single Root Input Output Virtualization (SR-IOV), multiple network interface cards may have the same MAC address. Consequently, the VM may fail to acquire an IP address from a

DHCP server and **NetworkManager-wait-online.service** may fail to start at boot time.

Workaround: Do not install the **initscripts-rename-device** package so that existing devices will not rename to existing device names.

Jira:RHEL-79783[1]

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

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

Legacy clients that do not support EMS or TLS 1.3 now cannot connect to FIPS servers running on RHEL 9 and 10. Similarly, RHEL 9 and 10 clients in FIPS mode cannot connect to servers that only support TLS 1.2 without EMS. This in practice means that these clients cannot connect to servers on RHEL 6, RHEL 7 and non-RHEL legacy operating systems. This is because the legacy 1.0.x versions of OpenSSL do not support EMS or TLS 1.3.

In addition, connecting from a FIPS-enabled RHEL client to a hypervisor such as VMWare ESX now fails with a **Provider routines::ems not enabled** error if the hypervisor uses TLS 1.2 without EMS. To work around this problem, update the hypervisor to support TLS 1.3 or TLS 1.2 with the EMS extension. For VMWare vSphere, this means version 8.0 or later.

For more information, see TLS Extension "Extended Master Secret" enforced with Red Hat Enterprise Linux 9.2 and later.

Jira:RHEL-13340[1]

10.19. RHEL IN CLOUD ENVIRONMENTS

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. There is currently no workaround for this issue.

Jira:RHELPLAN-120807[1]

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 sulogin(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.

Workaround: 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]

Enabling Hyper-V enlightenments in some cases does not improve CPU optimization

On virtual machines (VM) that use a Windows guest operating system, enabling Hyper-V enlightenments in some cases does not result in the expected improvement in the CPU usage of the VM. There is currently no workaround for this issue.

Jira:RHEL-17331[1]

Memory hot-plug possible on VMware when the memory size does not align with memory block size

Currently, it is possible to attempt hot-plugging memory to a RHEL 9 guest on VMware hypervisor even if the memory size of the attached memory dpes not align with the size of the individual memory blocks. However, attaching memory in this manner always fails with a **Block size unaligned hotplug range** error.

Workaround: Only hot-plug memory that is divisible by the configured memory block size on the guest. To obtain the memory block size, use the **lsmem** command. For further information, see The Red Hat KnowledgeBase.

Jira:RHEL-81748[1]

BIOS or UEFI supported Hyper-V Windows Server 2016 VM fails to boot if a host uses the AMD EPYC CPU processor

With the Hyper-V enabled setting, Hyper-V Windows Server 2016 VM fails to boot on the AMD EPYC CPU host.

Workaround: Check for the following log message:

```
kvm: Booting SMP Windows KVM VM with !XSAVES && XSAVEC.
If it fails to boot try disabling XSAVEC in the VM config.
```

And try adding **xsavec=off** to **-cpu cmdline** to boot Hyper-V Windows Server 2016 VM.

Jira:RHEL-38957[1]

kdump fails to complete on the Azure Confidential VMs

When you experience a kernel crash on a Red Hat Enterprise Linux VM on the Azure Confidential VM instances, in this case DCv5 and ECv5 series, the **kdump** process may not complete and the VM becomes unresponsive. As a result, after a forced reboot, there is a **vmcore-incomplete** file.

Jira:RHEL-70228[1]

10.20. SUPPORTABILITY

Timeout when running sos report on IBM Power Systems, Little Endian

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
```

Jira:RHELPLAN-51452[1]

10.21. CONTAINERS

UBI images are not reproducible

The **podman build** and **buildah build** commands avoid introducing inconsistencies between builds that use the same set of inputs when you invoke them with the following arguments:

- **--rewrite-timestamp**
- **--source-date-epoch**, an equivalent build argument or environment value that you set when starting the build.

To work around this problem, invoke the **podman build** or **buildah build** commands with the **--rewrite-timestamp** and **--source-date-epoch** arguments to minimize build inconsistencies. Additionally, update tools invoked in **RUN** instructions to avoid producing nondeterministic output when the **$SOURCE_DATE_EPOCH** environment variable is set.

Some tools or tool versions might still produce nondeterministic output, and you might not be able to build specific images reproducibly.

Jira:RHEL-62749

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.
```

Workaround: 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]

FIPS bootc image creation fails on FIPS enabled host

Building a disk image on a host by using Podman with enabled the FIPS mode fails with the exit code 3 because of the update-crypto-policies package:

```
# Enable the FIPS crypto policy
# crypto-policies-scripts is not installed by default in RHEL-10
RUN dnf install -y crypto-policies-scripts && update-crypto-policies --no-reload --set FIPS
```

Workaround: Build the bootc image with FIPS mode disabled.

Jira:RHELDOCS-19539

10.22. RHEL LIGHTSPEED

Command-line assistant configuration file changes are not applied immediately

When making changes in the **etc/xdg/command-line-assistant/config.toml** configuration file, it takes around 30 to 60 seconds for the command-line assistant daemon to recognize the changes, instead of applying the changes immediately. The command-line assistant is also missing the **reload** functionality.

Workaround: Follow the steps:

1. Make the changes that you need to the **config.toml** configuration file.
2. Run the following command:

```
# systemctl restart clad
```

Jira:RHELDOCS-19734[1]

CHAPTER 11. AVAILABLE BPF FEATURES

A complete list of the Berkeley Packet Filter (BPF) features that are available in this version of Red Hat Enterprise Linux 9 is provided in this chapter. 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 11.1. System configuration and other options

Option	Value
unprivileged_bpf_disabled	2 (bpf() syscall restricted to privileged users, admin can change)
bpf_jit_enable	1 (enabled)
bpf_jit_harden	1 (enabled)
bpf_jit_kallsyms	1 (enabled)
bpf_jit_limit	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 11.2. 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_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_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_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_strtol, 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_get_current_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_get_current_ancestor_cgroup_id, 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_get_current_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_get_current_ancestor_cgroup_id, 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
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_strtol, 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
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_get_current_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_skc_lookup_tcp, bpf_tcp_check_syncookie, bpf_strtol, bpf_strtoul, 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_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_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_strtol, 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_current_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_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_current_ancestor_cgroup_id, 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_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_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_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_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_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_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_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_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_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_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_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_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_get_current_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_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_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
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_strtol, 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_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_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_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_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_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_get_current_cgroup_id, 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_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
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_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_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_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_writable	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_strtol, 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_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_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	
struct_ops	
ext	
lsm	

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_get_current_cgroup_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_get_current_ancestor_cgroup_id, 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

Program type	Available helpers
netfilter	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_task, bpf_get_numa_node_id, 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_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

Table 11.3. 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

Map type	Available
devmap	yes
sockmap	yes
cpumap	yes
xskmap	yes
sockhash	yes
cgroup_storage	yes
reuseport_sockarray	yes
percpu_cgroup_storage	yes
queue	yes
stack	yes
sk_storage	yes
devmap_hash	yes
struct_ops	yes
ringbuf	yes
inode_storage	yes
task_storage	yes
bloom_filter	yes
user_ringbuf	yes
cgrp_storage	yes
arena_map	yes

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	Jira:RHEL-80163, Jira:RHEL-81141, Jira:RHEL-109034, Jira:RHEL-109892, Jira:RHEL-109889, Jira:RHEL-109885, Jira:RHEL-107585, Jira:RHEL-107005, Jira:RHEL-104593, Jira:RHEL-104591, Jira:RHEL-95444, Jira:RHEL-89753, Jira:RHEL-89745, Jira:RHEL-89736, Jira:RHEL-79673, Jira:RHEL-61347
ModemManager	Jira:RHEL-68732
NetworkManager	Jira:RHEL-85770, Jira:RHEL-85765, Jira:RHEL-83061, Jira:RHEL-24337, Jira:RHEL-5852, Jira:RHEL-24622, Jira:RHELPLAN-58745, Jira:RHEL-17619
NetworkManager-libreswan	Jira:RHEL-85768
Release Notes	Jira:RHELDOCS-20446, Jira:RHELDOCS-21013, Jira:RHELDOCS-18935, Jira:RHELDOCS-16861, Jira:RHELDOCS-17520, Jira:RHELDOCS-17803, Jira:RHELDOCS-19072, Jira:RHELDOCS-19635, Jira:RHELDOCS-20472, Jira:RHELDOCS-21350, Jira:RHELDOCS-19754, Jira:RHELDOCS-19889, Jira:RHELDOCS-20146, Jira:RHELDOCS-18158, Jira:RHELDOCS-17532, Jira:RHELDOCS-17508, Jira:RHELDOCS-19022, Jira:RHELDOCS-19284, Jira:RHELDOCS-18312, Jira:RHELDOCS-18480, Jira:RHELDOCS-19224, Jira:RHELDOCS-19028, Jira:RHELDOCS-19029, Jira:RHELDOCS-18959, Jira:RHELDOCS-19013, Jira:RHELDOCS-19012, Jira:RHELDOCS-19080, Jira:RHELDOCS-19050, Jira:RHELDOCS-19093, Jira:RHELDOCS-19149, Jira:RHELDOCS-19133, Jira:RHELDOCS-19171, Jira:RHELDOCS-19147, Jira:RHELDOCS-19139, Jira:RHELDOCS-19135, Jira:RHELDOCS-19137, Jira:RHELDOCS-19154, Jira:RHELDOCS-19143, Jira:RHELDOCS-19151, Jira:RHELDOCS-19115, Jira:RHELDOCS-18531, Jira:RHELDOCS-20135, Jira:RHELDOCS-16756, Jira:RHELDOCS-16612, Jira:RHELDOCS-17102, Jira:RHELDOCS-17309, Jira:RHELDOCS-17545, Jira:RHELDOCS-17518, Jira:RHELDOCS-17989, Jira:RHELDOCS-17702, Jira:RHELDOCS-17917, Jira:RHELDOCS-19193, Jira:RHELDOCS-19603, Jira:RHELDOCS-20196, Jira:RHELDOCS-16979, Jira:RHELDOCS-19846, Jira:RHELDOCS-20471
anaconda	Jira:RHEL-10216, Jira:RHEL-2250, Jira:RHEL-17205, Jira:RHEL-63237, Jira:RHELPLAN-168262, Jira:RHELPLAN-110940, Jira:RHELPLAN-53644, Jira:RHEL-4707, Jira:RHEL-4711, Jira:RHELPLAN-94811, Jira:RHEL-4741, Jira:RHELPLAN-130370, Jira:RHEL-4762, Jira:RHEL-4737, Jira:RHEL-9633, Jira:RHEL-14005, Jira:RHEL-80867, Jira:RHEL-82566, Jira:RHEL-78272, Jira:RHEL-13837, Jira:RHEL-66155

Component	Tickets
ansible-collection-microsoft-sql	Jira:RHEL-69311
ansible-pcp	Jira:RHEL-78306
azure-vm-utils	Jira:RHEL-88789
bacula	Jira:RHEL-6856
bind	Jira:RHELPLAN-90604
bootc	Jira:RHEL-90381
bootc-image-builder-container	Jira:RHEL-34807
buildah	Jira:RHEL-115166
ca-certificates	Jira:RHEL-54695
chrony	Jira:RHEL-95016
cockpit	Jira:RHEL-87397, Jira:RHEL-92062
cockpit-machines	Jira:RHEL-31993
cockpit-session-recording	Jira:RHEL-96905
container-tools	Jira:RHEL-69742, Jira:RHEL-67859
crash	Jira:RHEL-76270
createrepo_c	Jira:RHEL-67689, Jira:RHELPLAN-112860
crypto-policies	Jira:RHEL-91839, Jira:RHEL-104607, Jira:RHEL-103793, Jira:RHEL-2735, Jira:RHEL-112697
cups-filters	Jira:RHEL-6519
cyrus-sasl	Jira:RHELPLAN-94096

Component	Tickets
device-mapper-multipath	Jira:RHEL-78758, Jira:RHEL-82534, Jira:RHELPLAN-105944, Jira:RHELPLAN-99108, Jira:RHELPLAN-66975
distribution	Jira:RHEL-96056, Jira:RHEL-6973, Jira:RHEL-18157, Jira:RHEL-68141, Jira:RHEL-22385
dnf	Jira:RHEL-82310, Jira:RHEL-84512, Jira:RHEL-65817, Jira:RHEL-71125, Jira:RHELPLAN-118420
dnf-plugins-core	Jira:RHEL-100157
dyninst	Jira:RHEL-87002
edk2	Jira:RHELPLAN-69533, Jira:RHEL-82759, Jira:RHEL-68418
elfutils	Jira:RHEL-86971
fapolicyd	Jira:RHEL-63090, Jira:RHEL-69136, Jira:RHEL-21777, Jira:RHELPLAN-112355, Jira:RHEL-24345, Jira:RHEL-520
fence-agents	Jira:RHEL-79798, Jira:RHEL-68321, Jira:RHEL-7601, Jira:RHEL-88568, Jira:RHEL-82193, Jira:RHEL-13088
fips-provider-next	Jira:RHEL-105009, Jira:RHEL-96056
firewalld	Jira:RHEL-17708
gcc	Jira:RHEL-75806
gcc-toolset-15	Jira:RHEL-81741
gdb	Jira:RHEL-50069, Jira:RHEL-91381
gdm	Jira:RHEL-95837
gimp	Jira:RHEL-40106
glibc	Jira:RHEL-56627, Jira:RHEL-83017, Jira:RHEL-44920, Jira:RHEL-101986, Jira:RHEL-93320, Jira:RHEL-24168, Jira:RHEL-50086, Jira:RHEL-47403, Jira:RHEL-71922, Jira:RHEL-72017, Jira:RHEL-49549, Jira:RHEL-48820, Jira:RHEL-68805, Jira:RHEL-59712
gnome-control-center	Jira:RHEL-68152
gnome-settings-daemon	Jira:RHEL-11910

Component	Tickets
gnupg2	Jira:RHELPLAN-117566
gnutls	Jira:RHELPLAN-128129
go-rpm-macros	Jira:RHEL-7366
golang	Jira:RHELPLAN-129104, Jira:RHELPLAN-123778
gpsd-minimal	Jira:RHEL-90132
gtk3	Jira:RHEL-11924
httpd	Jira:RHEL-41069
initscripts	Jira:RHEL-79783
ipa	Jira:RHEL-84277, Jira:RHEL-30658, Jira:RHEL-67913, Jira:RHELPLAN-121751, Jira:RHELPLAN-113281, Jira:RHEL-12154, Jira:RHEL-4955
ipa-healthcheck	Jira:RHEL-4957
iproute	Jira:RHEL-90492, Jira:RHEL-94662
irqbalance	Jira:RHEL-89986
jmc-core	Jira:RHELPLAN-88788
kdump-anaconda-addon	Jira:RHEL-11196
kernel	Jira:RHELPLAN-102815, Jira:RHELPLAN-102321, Jira:RHELPLAN-108169, Jira:RHELPLAN-154595, Jira:RHELPLAN-153754, Jira:RHELPLAN-157294, Jira:RHELPLAN-147783, Jira:RHELPLAN-96004, Jira:RHELPLAN-99859, Jira:RHELPLAN-135779, Jira:RHELPLAN-114103, Jira:RHELPLAN-97394, Jira:RHELPLAN-134771, Jira:RHELPLAN-141042
kernel / Accelerators	Jira:RHEL-38583
kernel / Core	Jira:RHEL-25967
kernel / Crypto	Jira:RHEL-20145

Component	Tickets
kernel / Debugging-Tracing / Perf	Jira:RHEL-60216, Jira:RHEL-53585, Jira:RHEL-52654, Jira:RHEL-47454, Jira:RHEL-47444, Jira:RHEL-47424, Jira:RHEL-45095, Jira:RHEL-20110, Jira:RHEL-20094, Jira:RHEL-23496
kernel / Debugging-Tracing / rtla	Jira:RHEL-94317, Jira:RHEL-86051, Jira:RHEL-77358
kernel / Networking	Jira:RHEL-88890, Jira:RHEL-76845, Jira:RHEL-88552, Jira:RHEL-88551
kernel / Networking / IPSec	Jira:RHEL-30141, Jira:RHEL-1015
kernel / Networking / NIC Drivers	Jira:RHEL-73517, Jira:RHEL-63642, Jira:RHEL-9897, Jira:RHEL-36283
kernel / Networking / Netfilter	Jira:RHEL-81900
kernel / Platform Enablement	Jira:RHEL-2564
kernel / Platform Enablement / ppc64	Jira:RHEL-15404, Jira:RHEL-28702
kernel / RDMA	Jira:RHEL-86016
kernel / Security / Other	Jira:RHEL-15599
kernel / Storage / Multiple Devices (MD)	Jira:RHEL-30730
kernel / Storage / Storage Drivers	Jira:RHEL-9301, Jira:RHEL-8171, Jira:RHEL-8164, Jira:RHEL-8466, Jira:RHEL-8104, Jira:RHEL-25730
kernel / Virtualization	Jira:RHEL-1138
kernel / Virtualization / Hyper-V	Jira:RHEL-70228, Jira:RHEL-29919
kernel / Virtualization / KVM	Jira:RHEL-11431, Jira:RHEL-10019, Jira:RHEL-7212, Jira:RHEL-17331, Jira:RHEL-45585, Jira:RHEL-32892, Jira:RHEL-38957
kernel / Virtualization / Public Cloud Enablement	Jira:RHEL-81748

Component	Tickets
kernel-rt	Jira:RHELPLAN-153123
kernel-rt / Other	Jira:RHEL-76757, Jira:RHEL-9318
kexec-tools	Jira:RHEL-33413, Jira:RHELPLAN-129876, Jira:RHEL-11471, Jira:RHELPLAN-115732
keylime	Jira:RHEL-78418, Jira:RHEL-11867, Jira:RHEL-1518
kmod	Jira:RHELPLAN-126922
kmod-kvdo	Jira:RHEL-8354
kpatch	Jira:RHEL-85579
krb5	Jira:RHELPLAN-114497, Jira:RHEL-4875, Jira:RHEL-4888
libabigail	Jira:RHEL-16629
libdnf	Jira:RHELPLAN-128381
libotr	Jira:RHELPLAN-122108
libvirt	Jira:RHELPLAN-139536, Jira:RHELPLAN-119912
libvirt / General	Jira:RHEL-72976, Jira:RHEL-11435, Jira:RHEL-7043, Jira:RHEL-89415
libxcrypt	Jira:RHELPLAN-106338
llvm	Jira:RHEL-81006
lorax-templates-rhel	Jira:RHEL-91930
lvm2	Jira:RHEL-67039, Jira:RHELPLAN-107107
mysql	Jira:RHELPLAN-92864
nfs-utils	Jira:RHELPLAN-120807
nmstate	Jira:RHEL-88993, Jira:RHEL-85784, Jira:RHEL-80725, Jira:RHEL-80418, Jira:RHEL-67631, Jira:RHEL-32495
nodejs	Jira:RHEL-35990, Jira:RHEL-90821

Component	Tickets
nss	Jira:RHEL-103366, Jira:RHEL-127671
nvme-stas	Jira:RHELPLAN-58357
open-vm-tools	Jira:RHELPLAN-106947
opencryptoki	Jira:RHEL-73344
opensc	Jira:RHEL-96029
openscap	Jira:RHELPLAN-145263
openslp	Jira:RHEL-6995
openssh	Jira:RHEL-104580, Jira:RHELPLAN-113842, Jira:RHEL-45727
openssl	Jira:RHEL-80854, Jira:RHEL-90854, Jira:RHEL-95239, Jira:RHELPLAN-148207, Jira:RHELPLAN-50959, Jira:RHELPLAN-48241, Jira:RHEL-40605, Jira:RHELPLAN-113856, Jira:RHELPLAN-139207
osbuild-composer	Jira:RHEL-4649
oscap-anaconda-addon	Jira:RHEL-1824, Jira:RHELPLAN-44202
pacemaker	Jira:RHEL-86143, Jira:RHEL-84018
pam	Jira:RHEL-15324
pause-container	Jira:RHELPLAN-127619
pcp	Jira:RHEL-83154
pcs	Jira:RHEL-76177, Jira:RHEL-76170, Jira:RHEL-76154, Jira:RHEL-76153, Jira:RHEL-76060, Jira:RHEL-76059, Jira:RHEL-76055, Jira:RHEL-35420, Jira:RHEL-92044, Jira:RHEL-34781
perl-DBD-MySQL	Jira:RHEL-77083
pkcs11-provider	Jira:RHEL-105625
pki-core	Jira:RHEL-98719, Jira:RHELPLAN-121754

Component	Tickets
podman	Jira:RHEL-88521, Jira:RHEL-88472, Jira:RHEL-88464, Jira:RHEL-88307, Jira:RHEL-110317, Jira:RHEL-32267, Jira:RHEL-70217, Jira:RHEL-88121, Jira:RHELPLAN-117005
postgis	Jira:RHEL-81603
procps-ng	Jira:RHEL-46760
python-blivet	Jira:RHEL-8008, Jira:RHEL-8012
python-drgn	Jira:RHEL-86264
python3.11-lxml	Jira:RHELPLAN-143480
qemu-kvm	Jira:RHEL-86032, Jira:RHEL-57677, Jira:RHELPLAN-81033, Jira:RHELPLAN-75969, Jira:RHELPLAN-114513, Jira:RHELPLAN-99854, Jira:RHELPLAN-63771, Jira:RHELPLAN-150884, Jira:RHELPLAN-118495, Jira:RHEL-7478, Jira:RHEL-62742, Jira:RHEL-67699, Jira:RHEL-66229
qemu-kvm / Devices	Jira:RHEL-1220
qemu-kvm / Devices / CPU Models	Jira:RHEL-17614
qemu-kvm / Graphics	Jira:RHEL-7135
qemu-kvm / Live Migration	Jira:RHEL-7096
qemu-kvm / Networking	Jira:RHEL-7337, Jira:RHEL-7335, Jira:RHEL-7336, Jira:RHEL-333, Jira:RHEL-21867
qemu-kvm / Storage	Jira:RHEL-82906
rear	Jira:RHEL-56045
redhat-release	Jira:RHEL-86164
resource-agents	Jira:RHEL-88429, Jira:RHEL-88035, Jira:RHEL-85220, Jira:RHEL-7688, Jira:RHEL-32265
restore	Jira:RHELPLAN-94704
rhel-bootc-container	Jira:RHEL-33208

Component	Tickets
rhel-system-roles	Jira:RHEL-99089, Jira:RHEL-95874, Jira:RHEL-104659, Jira:RHEL-102637, Jira:RHEL-84930, Jira:RHEL-104676, Jira:RHEL-104891, Jira:RHEL-95885, Jira:RHEL-82825, Jira:RHEL-103889, Jira:RHEL-88299, Jira:RHEL-84951, Jira:RHEL-101678, Jira:RHEL-87579, Jira:RHEL-17564, Jira:RHEL-107049, Jira:RHEL-107015, Jira:RHEL-106733, Jira:RHEL-105095, Jira:RHEL-103575, Jira:RHEL-101663, Jira:RHEL-94444, Jira:RHEL-93296, Jira:RHEL-88772, Jira:RHEL-88314, Jira:RHEL-88251, Jira:RHEL-88241, Jira:RHEL-85872, Jira:RHEL-85702, Jira:RHEL-84940, Jira:RHEL-84920, Jira:RHEL-84910, Jira:RHEL-84362, Jira:RHEL-85079, Jira:RHEL-81755, Jira:RHELPLAN-95747, Jira:RHELPLAN-133165, Jira:RHEL-1172
rng-tools	Jira:RHEL-91119
rpm	Jira:RHEL-35619, Jira:RHEL-52772
rsyslog	Jira:RHEL-66274, Jira:RHEL-92262
rteval	Jira:RHEL-97540
rust	Jira:RHEL-81601
rust-rpm-sequoia	Jira:RHEL-126412, Jira:RHEL-111478
s390utils	Jira:RHEL-73342
samba	Jira:RHEL-89873
scap-security-guide	Jira:RHEL-111009, Jira:RHEL-1800, Jira:RHELPLAN-107318
seabios	Jira:RHEL-7110
selinux-policy	Jira:RHEL-87744, Jira:RHEL-82674, Jira:RHEL-69526, Jira:RHEL-95342, Jira:RHEL-11792, Jira:RHELPLAN-115609, Jira:RHEL-28814
shim	Jira:RHEL-18969
sos	Jira:RHEL-71825, Jira:RHEL-62972, Jira:RHEL-67097, Jira:RHEL-73028, Jira:RHEL-81187, Jira:RHEL-81634, Jira:RHEL-84078, Jira:RHELPLAN-51452
sssd	Jira:RHELPLAN-44204
stalld	Jira:RHEL-108827

Component	Tickets
stunnel	Jira:RHEL-52317
subscription-manager	Jira:RHEL-84890, Jira:RHEL-29178, Jira:RHELPLAN-146101, Jira:RHELPLAN-137234
subversion	Jira:RHEL-79948
sysstat	Jira:RHEL-12009, Jira:RHEL-26275
systemd	Jira:RHEL-50534, Jira:RHELPLAN-100926, Jira:RHEL-6105, Jira:RHEL-92781
systemtap	Jira:RHEL-87000
tftp	Jira:RHEL-77491
tigervnc	Jira:RHELPLAN-114314
toolbox-container	Jira:RHEL-84787
trustee-guest-components	Jira:RHEL-68141
tuned	Jira:RHELPLAN-129881, Jira:RHEL-79914
tzdata	Jira:RHEL-105043
ubi9-container	Jira:RHEL-62749
unbound	Jira:RHELPLAN-117492
valgrind	Jira:RHEL-86998, Jira:RHEL-75468
valkey	Jira:RHEL-89978
vdo	Jira:RHEL-83857, Jira:RHEL-30525
virt-v2v	Jira:RHELPLAN-147926, Jira:RHEL-13340
virtio-win	Jira:RHEL-11810, Jira:RHEL-11366, Jira:RHEL-1609, Jira:RHEL-869
virtio-win / distribution	Jira:RHEL-1860, Jira:RHEL-574
virtio-win / virtio-win-prewhql	Jira:RHEL-1084, Jira:RHEL-935, Jira:RHEL-12118, Jira:RHEL-1212, Jira:RHEL-53962

Component	Tickets
virtiofsd	Jira:RHEL-87161
webkit2gtk3	Jira:RHEL-4157
xdp-tools	Jira:RHEL-3382
other	Jira:RHELDOCS-20446, Jira:RHELDOCS-20639, Jira:RHELDOCS-20303, Jira:RHELDOCS-21029, Jira:RHELDOCS-19500, Jira:RHELDOCS-21016, Jira:RHELDOCS-21025, Jira:RHELDOCS-21026, Jira:RHELDOCS-20633, Jira:RHELDOCS-20546, Jira:RHELDOCS-20421, Jira:RHELDOCS-21241, Jira:RHELDOCS-21230, Jira:RHELDOCS-21218, Jira:RHELDOCS-21313, Jira:RHELDOCS-17733, Jira:RHELDOCS-20781, Jira:RHELDOCS-21013, Jira:RHELDOCS-21314, Jira:RHELDOCS-21315, Jira:RHELDOCS-21316, Jira:RHELDOCS-21317, Jira:RHELDOCS-21318, Jira:RHELDOCS-21319, Jira:RHELDOCS-21320, Jira:RHELDOCS-21322, Jira:RHELDOCS-17040, Jira:RHEL-88550, Jira:RHELDOCS-20059, Jira:RHELPLAN-27394, Jira:RHELPLAN-27737, Jira:RHELDOCS-18935, Jira:RHELDOCS-21376, Jira:RHELDOCS-20258, Jira:RHELDOCS-16861, Jira:RHELDOCS-17520, Jira:RHELDOCS-17752, Jira:RHELDOCS-17803, Jira:RHELDOCS-17468, Jira:RHELDOCS-19664, Jira:RHELDOCS-19635, Jira:RHELDOCS-21350, Jira:RHELDOCS-19754, Jira:RHELDOCS-19455, Jira:RHEL-91106, Jira:RHELDOCS-19815, Jira:RHELDOCS-19716, Jira:RHELDOCS-19809, Jira:RHELDOCS-19810, Jira:RHELDOCS-19523, Jira:RHELDOCS-20146, Jira:RHELDOCS-20283, Jira:RHELDOCS-20464, Jira:RHELDOCS-20519, Jira:RHELDOCS-20718, Jira:RHELDOCS-17532, Jira:RHELDOCS-17508, Jira:RHELDOCS-19004, Jira:RHELDOCS-18312, Jira:RHELDOCS-18480, Jira:RHELDOCS-18701, Jira:RHELDOCS-18702, Jira:RHELDOCS-18703, Jira:RHELDOCS-19224, Jira:RHELDOCS-19028, Jira:RHELDOCS-19029, Jira:RHELDOCS-18592, Jira:RHELDOCS-18593, Jira:RHELDOCS-18803, Jira:RHELDOCS-18207, Jira:RHELDOCS-19013, Jira:RHELDOCS-19021, Jira:RHELDOCS-19012, Jira:RHELDOCS-19068, Jira:RHELDOCS-19069, Jira:RHELDOCS-19080, Jira:RHELDOCS-19050, Jira:RHELDOCS-19093, Jira:RHELDOCS-19115, Jira:RHELPLAN-67314, Jira:RHELPLAN-110763, Jira:RHELPLAN-69554, Jira:RHELPLAN-113995, Jira:RHELDOCS-20135, Jira:RHELPLAN-122745, Jira:RHELPLAN-99136, Jira:RHELPLAN-103232, Jira:RHELPLAN-60153, Jira:RHELPLAN-88246, Jira:RHELPLAN-100087, Jira:RHELPLAN-100639, Jira:RHELPLAN-113659, 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:RHELDOCS-17015, Jira:RHELDOCS-18049, Jira:RHELDOCS-17135, Jira:RHELDOCS-17545, Jira:RHELDOCS-17038, Jira:RHELDOCS-17495, Jira:RHELDOCS-17518, Jira:RHELDOCS-17462, Jira:RHELDOCS-18106, Jira:RHELDOCS-17782, Jira:RHELDOCS-19193, Jira:RHELPLAN-157225, Jira:RHELPLAN-10061, Jira:RHELPLAN-110191, Jira:RHELPLAN-96940, Jira:RHELPLAN-117234, Jira:RHELPLAN-119001, Jira:RHELPLAN-119852, Jira:RHELPLAN-119655,

Component	Tickets
	Jira:RHELPLAN-112043, Jira:RHELPLAN-121205, Jira:RHELPLAN-121043, Jira:RHELPLAN-109613, Jira:RHELPLAN-145001, Jira:RHELDOCS-19603, Jira:RHELDOCS-18064, Jira:RHELDOCS-16427, Jira:RHELPLAN-150080, Jira:RHELPLAN-154195, Jira:RHELPLAN-83423, Jira:RHELDOCS-17719, Jira:RHELDOCS-19945, Jira:RHELDOCS-18720, Jira:RHELDOCS-18435, Jira:RHELDOCS-18863, Jira:RHELDOCS-19728, Jira:RHELDOCS-19539, Jira:RHELDOCS-19734, Jira:RHELDOCS-19948, Jira:RHELDOCS-19496

APPENDIX B. REVISION HISTORY

0.0-0

Wed 12 Nov 2025, Valentina Ashirova (vaashiro@redhat.com)

- Release of the Red Hat Enterprise Linux 9.7 Release Notes.